



INSTITUTE FOR DEFENSE ANALYSES

## **Analyzing Adversaries as Complex Adaptive Systems**

Dale E. Lichtblau, Task Leader

Brian A. Haugh  
Gregory N. Larsen  
Terry Mayfield

October 2006

Approved for public release;  
unlimited distribution.

IDA Paper P-3868

Log: H 04-000048

**This work was conducted under IDA's independent research program. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.**

**© 2004, 2006 Institute for Defense Analyses, 4850 Mark Center Drive, Alexandria, Virginia 22311-1882 • (703) 845-2000.**

**This material may be reproduced by or for the U.S. Government.**

INSTITUTE FOR DEFENSE ANALYSES

IDA Paper P-3868

## **Analyzing Adversaries as Complex Adaptive Systems**

Dale E. Lichtblau, Task Leader

Brian A. Haugh  
Gregory N. Larsen  
Terry Mayfield



## **Preface**

---

The work described in this paper was sponsored by the Institute for Defense Analyses (IDA) under an internal Central Research Project (CRP) and reports on research largely conducted during the period January 2002–September 2003. The task was an investigation of utility of treating terrorist groups as complex adaptive systems and thus subject to analysis using agent-based modeling technology. The effort was supported by Mr. Adam Stockton, at the time a rising junior at Yale University, who worked as an intern during the summer of 2003. The effort was supported by Mr. Adam Stockton, a rising junior at Yale University, who worked as an intern during the summer of 2003.

Mr. Mathew MacArthur also provided invaluable support to this research effort. Moreover, he provided an excellent constructive critique of the executable model developed during the course of this project.

This document was reviewed by IDA Fellow, Dr. Richard J. Ivanetich, and IDA Research Staff Member, Dr. L. Roger Mason, Jr.



## Contents

---

Executive Summary.....	ES-1
1. Introduction.....	1
1.1 Background.....	2
1.2 Statement of Problem.....	3
1.3 Objectives of the Research.....	4
1.4 Technical Approach.....	5
1.5 Future Applications.....	7
1.6 Definitions.....	8
1.6.1 Terrorism.....	8
1.6.2 Complex Adaptive Systems.....	12
1.6.3 Agent-Based Modeling.....	13
1.6.4 Complexity Theory.....	17
1.6.5 Social Network Analysis.....	21
1.7 Organization of the Report.....	23
2. Terrorist Organizations Viewed as Complex Adaptive Systems.....	25
2.1 Basic Idea of the Model.....	25
2.2 Basic Model Parameters.....	26
2.2.1 Simulation Controls.....	27
2.2.1.1 Environment.....	28
2.2.1.2 Output.....	29
2.2.2 Actors Rules.....	31
2.2.2.1 Citizens.....	32
2.2.2.2 Police.....	33
2.2.2.3 Terrorists.....	34

2.2.3 Advanced Options.....	35
2.2.3.1 Sympathy Constraints.....	36
2.2.3.2 Random Number Options.....	38
2.3 Actor Behavior.....	39
2.3.1 The Sympathy Continuum.....	39
2.3.2 Alliances.....	40
2.3.3 Adaptation.....	40
2.3.4 Innovation.....	41
2.3.4.1 Use of IKB.....	42
2.3.4.2 Use of Open Cycle.....	44
2.4 Design of Experiments.....	45
3. Preliminary Results and Interpretations.....	48
4. Conclusions.....	62
4.1 Tactical value of Agent-Based Modeling.....	62
4.1.1 The Need to Validate Agent-Based Models.....	63
4.2 Strategic Value of Agent-Based Modeling/Simulation Research.....	64
4.3 The Hidden Costs of Agent-Based Modeling.....	66
4.4 The Ultimate Value of Agent-Based Modeling.....	67
Appendix A. Terrorism CRPMt.....	A-1
Appendix B. An Analysis of Some Key Assumptions of the Terrorism Model by Matthew C. MacArthur.....	B-1
References.....	Ref-1
Acronyms and Abbreviations.....	Acros-1



## Figures

---

Figure 1. Total International Terrorist Attacks, 1981 – 2002.....	20
Figure 2. Detrended Data .....	20
Figure 3. Phase Plot of the Detrended Data.....	21
Figure 4. Initial Window of the Graphical User Interface .....	27
Figure 5. Landscape as Torus .....	28
Figure 6. Output File Example .....	29
Figure 7. Viewer Display.....	30
Figure 8. Actor Rules Setup Window .....	32
Figure 9. Police Precision.....	34
Figure 10. Advanced Options .....	36
Figure 11. The Sympathy Continuum .....	40
Figure 12. Design of Experiments Interface Window .....	46
Figure 13. Design of Experiments.....	47
Figure 14. Number of Police and Terrorists as Function of Attack Interval and Terrorist Links ..	51
Figure 15. Terrorist Events, Sympathy, and Competence as Function of Minimum Attack Interval and Maximum Terrorist Links .....	52
Figure 16. Number of Police, Terrorists, and Latent Terrorists as Function of Terrorist Attack Magnitude and Police Precision .....	53
Figure 17. Terrorist Events, Sympathy, and Competence as Function of Terrorist Attack Magnitude and Police Precision .....	54
Figure 18. Number of Police and Terrorists as Function of Terrorist Attack Magnitude and Police Precision of Three with Learning Enabled and Disabled .....	55
Figure 19. Terrorist Events, Sympathy, and Competence as Function of Terrorist Attack Magnitude and Police Precision of Three with Learning Enabled and Disabled ...	56
Figure 20. Number of Police, Terrorists, and Latent Terrorist as Function of Terrorist Attack Magnitude and Police Precision of Five with Learning Enabled and Disabled .....	57
Figure 21. Events, Sympathy, and Competence as Function of Terrorist Attack Magnitude and Police Precision of Five with Learning Enabled and Disabled .....	58
Figure 22. Number of Police, Terrorists, and Latent Terrorists as Function of Terrorist Magnitude and Maximum Terrorist Links with Learning Enabled and Disabled.....	59
Figure 23. Number of Events, Sympathy, and Competence as Function of Terrorist Magnitude and Maximum Terrorist Links with Learning Enabled and Disabled.....	60

Figure 24. Number of Police and Terrorists as Function of Terrorist Sympathy and Police Sympathy .....	61
Figure 25. Events, Sympathy, and Competence as Function of Terrorist Sympathy and Police Sympathy .....	62

## Tables

---

Table 1. Map of Simulation Result Graphs .....	50
--	----



## Executive Summary

---

The events on 11 Sep[tember] 2001 were a tragic, but decisive, reminder of the emergence of a formidable new kind of "enemy" in the world; an enemy that is widely dispersed, decentralized and whose many destructive parts are autonomous, mobile, and highly adaptive. The need for developing new complex systems theory inspired analytical tools and models for understanding the dynamics of this threat (and for providing insights into how to combat it) has never been greater. *If ever there was a time for complexity theory to come into its own within the military operations research community (much as mathematical search theory did in WW II), that time is now!*<sup>1</sup>

This quote by Mr. Andy Illyachinsk—a well-known and respected researcher at the Center for Naval Analyses (CNA)—captures precisely both the motivation behind and the basic question that underlies the efforts documented in this report. The motivation is to find additional ways to fight global terrorism. The basic question is, Are asymmetric threats—in particular, terrorism—usefully analyzable as complex adaptive systems? Put another way and more broadly, how relevant is complexity theory to the analysis of asymmetric threats?

Our overall objective can be summarized as the assessment of the following “argument.” The groups or organizations that pose asymmetric threats to the United States are complex adaptive systems (CASs). Complexity theory is that discipline devoted to the development of the theory and tools to analyze and come to better understand complex systems (whether adaptive or not). Agent-based modeling (ABM) has proved to be an especially well-suited tool for the systematic analysis of complex systems. Therefore, there’s *prima facie* evidence that complexity theory in general and ABM in particular are probably good methods—maybe even the best method—for analyzing terrorism and similar phenomena.

We argue that while terrorist groups considered *qua systems* are no doubt *adaptive* it is not obvious that they are *complex* in the formal or theoretical sense of that term. There-

---

<sup>1</sup> [http://www.cna.org/isaac/terrorism\\_and\\_cas.htm](http://www.cna.org/isaac/terrorism_and_cas.htm), February 6, 2002.

fore, it's not obvious that such groups—and the asymmetric threats they pose—are only amenable to analysis with the emerging techniques of complexity theory.

But even under the assumption that terrorist groups *are* complex systems, it is still arguable that ABM is an effective or even a viable, let alone the best, means to explore their complex behavior and to derive predictable—and therefore usable—results. The behavior under examination, the highly complicated behavior of human beings acting within a highly complex and dynamic world, is simply too complicated and too poorly understood to be modeled in anything but a trivial, simplistic, unrealistic, and completely enlightening way *within the time frames necessary to have tactical import*. The agent-based models that purport to simulate real-world social behavior are arguably “toys.” The process leading to their development—specifically, the process of making explicit the assumptions or premises regarding the fundamental factors governing the behavior of the agents in the model—are not without value, and potentially of inestimable value. But care must be exercised so as not to go beyond the endorsement of the assumptions to the belief that the system effects that emerge from the repeated execution of those assumptions is anything but an artifact of the model and rarely, if ever, a projection from or a mirror of the real world.

There is considerable value in complexity theory and ABM and social network analysis (SNA) and systems dynamics. The problems to which these analytical techniques are being applied certainly include the kinds of problems for which these tools were devised and which offer, moreover, the promise of eventual success. Real success, however, is not imminent. Complexity science is not verged upon a *coming into its own* within military operations research community, certainly for *tactical* advantage. The value of complexity theory vis-à-vis the Department of Defense (DoD), in our view, is currently *strategic*. It focuses on trying to understand the dynamics and underlying law-like rules of those aspects of systems behavior that are clearly governing the threats we face. For no other reason, it deserves our attention, our respect, and continued DoD interest. Tactical utility, except to the extent that greater appreciation and insight into complex systems behavior begins to seep into current tactical-oriented analyses, is not in immediate or even near-term offing.

Alas, this pessimistic conclusion is not argued for very cogently. Our research began optimistically with the hope that we could coax ABM techniques into lines of inquiring that overcame some of our persistent doubts about the soundness of many attempts to apply ABM (and complexity science) to national defense issues. We did not focus on a critique of the ABM approach so much as try to learn the science and apply it as best we could to a domain with which we are reasonably familiar. The *critique* of the ABM methodology is admittedly weak. The critique from the positive perspective is more cogent, but certainly not impregnable. It's easily imaginable that the trails we began to blaze with this research can be extended and eventually meet with more possible results. Accordingly, we encourage future researches to address the following issues when applying complexity science—broadly construed—to asymmetric threats:

- Do the asymmetric threats exhibit genuine complex behavior?
- If so, is it possible to predict, using ABM techniques, and threat behavior countered within an “effective response cycle”<sup>2</sup>?
- Even if symmetric threats are not the result of genuinely complex system behavior, might ABM and related techniques still have an important role to play in their analysis?
- Can more effective and easier to use modeling environments be developed?

---

<sup>2</sup> An “effective response cycle” is the time interval in which it is possible to mount a potentially effective response to an attack which initiates the time interval.





## 1. Introduction

---

This paper reports on a “voyage”—with apologies to Mr. R. H. Dana, Jr.—of roughly two years duration exploring the notion of *complex adaptive systems* and its value as an explanatory concept in strengthening U.S. national security. This particular excursion began in 2001 with (what we thought was) a well circumscribed focus: what value might ensue from treating asymmetric threats—specifically, transnational terrorism—as complex adaptive systems? If we were to assume that (transnational) terrorist networks were essentially “complex adaptive systems,” would modern computational modeling and simulation technologies that purported to model (and simulate) such systems be able to be marshaled, in some way, in an attack on (or defense from) such “asymmetric threats.”

This “voyage” took us too many “ports of call.” We perused the literature on terrorism, both ancient and modern. We reviewed the work of many contemporary researchers also engaged in studying terrorist behavior. We looked at “modern systems theory” and delved into the modern but abstruse notion of “complexity.” The concepts of “adaptation” and “learning” forced us to review the principles of Darwinian evolution and neo-evolutionary thinking. We looked at a broad range of contemporary computer simulation tools and environments. The present research was also informed by two previous Central Research Projects (CRPs)<sup>3</sup> and two direct Institute for Defense Analyses (IDA) tasks<sup>4</sup>.

The broad and diffuse nature of our work required us to decide between two forms of a final research report: a short, sharp, and well focused report of the findings and conclusions of a particular (well-focused, but limited) “experiment,” or a larger document that captured the broader tenor of the two year (or so) effort. We opted for the latter, but we use the “focused experiment” as the unifying thread of the report. What this means for the

---

<sup>3</sup> C5001 (Complexity Studies) and C5051 (Complexity Science).

<sup>4</sup> Task Order BB-5-1897 (Special Operations) and Task Order AJ-5-1944 (Complexity Science).

reader is that a lot of material is presented that, while relevant to the overall objective of the research, it may not be *directly* relevant to the central question: Can asymmetric threats be analyzed as complex adaptive systems?

This section provides a concise introduction to the problem being addressed, the objectives of the research, the technical approach employed, and possible future applications of the findings. It also provides a background context to the present research efforts.

## 1.1 Background

Most of the specific research project reported on in this paper is the most recent of several related projects within the Information Technology and Systems Division (ITSD)<sup>5</sup>. Three of these projects were funded by IDA’s Central Research Project (CRP) program for the purpose of enabling division researcher staff members to keep abreast of developments in the field known as complexity science and its possible applicability to national security and defense issues. Earlier efforts were principally “information collecting” focused. We read the current literature, attended relevant conferences<sup>6</sup>, and “networked” with the community of “like-minded” researchers. And while we examined many of the tools being used in complexity science research—mainly agent-based models—, we did not attempt in these earlier efforts to actually build such tools (i.e., agent-based models) on our own. As a consequence, we could not offer an assessment of the value of complexity science to DoD based on an examination of all aspects of the discipline. We lacked practical hands-on experience with the computer code that was being used to implement agent-based models for the analysis of complex systems. The present research effort addresses these earlier omissions. While informed with considerable and

---

<sup>5</sup> Formerly the Computer and Software Engineering Division (CSED).

<sup>6</sup> Including the following: *Preserving National Security in a Complex World: A Colloquium on Innovative Applications of Complexity Science in the Military, Government, and Civilian Industry*, Cambridge, MA, 12 – 14 September 1999; *Complexity: An Important New Framework for Understanding Warfare? A Symposium*, 28 February 2001, Center for Naval Analyses, Alexandria, VA; *Beyond Pont Estimates*, the 4<sup>th</sup> International Military Applications Society Conference, 21 – 23 May 2001, Quantico, VA; *High Altitude Thinking – International InfoMesa Summit*, 27 – 30 August 2001; *Project O’Bannon—Terrorist Networks: An Analysis*, 22 – 23 July, 2002, Quantico, VA; *Swarming: Network Enabled C4ISR*, 13-14 January 2003, McLean, VA.

invaluable knowledge of complexity science gained from our previous CRP work, the present effort focused on building and analyzing an executable (agent-based) model of a complex system. Accordingly, this paper focuses on this later effort, but two themes are interwoven throughout this report: the lessons learned from actually attempting to apply complexity science techniques to defense issues (the focus of the current effort), and a more general assessment of the utility of complexity science, based on both current and previous CRP work. It should also be noted that our views on the value of agent-based modeling were also informed by two direct (sponsor funded) tasks (both cited earlier). One was a major project sponsored by the Office of the Assistant Secretary for Science and Technology, Special Projects, conducted in 2001; it looked at the potential role of “complexity science” in DoD. The other was sponsored by the Office of the Assistant Secretary of Defense for Special Operations and Low Intensity Conflicts (SO/LIC) and focused on the use of multi-agent-based systems to support psychological operations (PSYOP) analysts.

## **1.2 Statement of the Problem**

Asymmetric adversaries (in particular, transnational terrorist networks) seem to exhibit complex adaptive system (CAS) behavior. The inherent complexity of such systems make them especially resistant to common analytical techniques. Moreover, terrorist networks (and any other organization engaged in illicit behavior) act covertly and, almost by definition, cannot be studied directly from outside of the organization.<sup>7</sup> There are, however, a number of indirect techniques used widely and successfully in the natural sciences. Modeling (and simulation)—a form of hypothesis framing and testing—is the paradigmatic scientific method for studying the behavior of systems for which direct examination is impractical if not theoretically impossible. The immediate problem is to determine if

---

<sup>7</sup> Sun-Ki Chai writes: “One obviously can not carry out a telephone or door-to-door survey among terrorists or guerrillas. Even where one can get access to them and ask questions, there are obvious problems of response bias. It is unlikely that active guerrillas will admit that their motives are anything but pure and entirely political; defecting or captured guerrillas have a tendency to say what they think their interrogators want to hear” (Chai 1993, pp. 100-101).

there is such a technique that can be used to advantage in studying complex adaptive systems, in particular, terrorist networks and similar asymmetric threats.

### 1.3 Objectives of the Research

The principal objective of this research is to determine—or at least to raise the issue in a methodical way—if agent-based modeling (ABM), social network analysis (SNA) and similar “systems” analysis tools and techniques are well suited for the study of asymmetric adversaries when viewed as CASs. The motivation for this research is four fold: one particular asymmetric threat—terrorism—has become a serious worldwide danger; ABMs and related computing-based technologies are being touted as the only viable means of investigating in a quantitative, and therefore, testable way these complex adaptive systems; and large amounts of taxpayer money are being expended in the hope of using these technologies to counter these growing threats to United States national security.

From another perspective, we aim to better define the benefits that are frequently alleged to derive from using these computationally intensive but now readily affordable techniques for analyzing complex systems that adapt to changing circumstances. We hope to provide some answers to the question, What exactly do these technologies have to offer?

Third, we hope that some practical and actionable insights can be gained from our work. For instance, we would like to see some predictions of CAS behavior that could be tested against the historical record. If a simulation results in behavior  $\mathfrak{B}$  from a set of initial conditions  $\mathcal{C}$  in terms of the rules of behavior embodied in a model  $\mathfrak{M}$ , and if we think some physical world<sup>8</sup> organization  $M$  behaves in accordance with the rules of  $\mathfrak{M}$ , then we would like to see physical world conditions  $C$  that are comparable to the initial model conditions  $\mathcal{C}$  and subsequent physical world behavior  $B$  similar to  $\mathfrak{B}$ . We could conclude that the rules of behavior that govern  $\mathfrak{M}$  are similar in essence to the rules of behavior

---

<sup>8</sup> The *physical* world is that world being modeled via a *conceptual* world. The physical world and the conceptual worlds jointly comprise the *real* world.

that govern M in the physical world. Knowledge of the latter would then provide a significant advantage in countering the threats posed by M.<sup>9</sup>

A final objective was to enable a small cadre of IDA researchers to stay abreast of developments in the field of complexity science, social network analysis, and system dynamics, especially as those techniques and technologies can be applied to support national security.

#### **1.4 Technical Approach**

Our technical approach to this research effort consisted of four stages. The first stage was a survey of the literature and information technology related to the problem. We consulted the relevant literature, both that which addressed terrorism as a political, cultural, or social phenomenon—historical and contemporary—and that which focused on theories of the various kinds of *system* behavior we thought to be exhibited by asymmetric adversaries. With respect to the latter we looked at complex adaptive systems, social network analysis theory, and system dynamics. We also surveyed the complexity and agent-based modeling literature, attended relevant conferences, and, as already noted, drew on considerable experience from previous CRP and DoD-sponsored tasks.

Since our objective was to assess the value of agent-based modeling technology for this particular problem, we also surveyed the technology arena, looking at various ABM development environments with an eye to adopting at least one platform to use in the actual building of a model and simulation to test our ideas. We looked at Swarm,<sup>10</sup> ABIR (Agent Based Identity Repertoire)<sup>11</sup>, Isaac/Einstein<sup>12</sup>, Ascape<sup>13</sup>, Project Albert<sup>14</sup>, and

---

<sup>9</sup>Note that achieving this latter goal will attest to the value of these technologies vis-à-vis certain asymmetric threats. The failure to come up with viable ideas to thwart terrorist threats using ABM and related technologies will not prove that these methods are of no value, however. One can always argue that we have just not been clever enough to find the right application of the technology to the problem at hand.

<sup>10</sup> <http://www.swarm.org/>

<sup>11</sup> [http://www.psych.upenn.edu/sacsec/abir/\\_private/Projects.htm](http://www.psych.upenn.edu/sacsec/abir/_private/Projects.htm)

<sup>12</sup> <http://www.cna.org/isaac/>

MANA (Map Aware Non-uniform Automata)<sup>15</sup> in varying degrees of detail. (Ackerman 2000 provides a comprehensive survey of researchers interested in modeling terrorist behavior.) Ultimately we decided to collaborate with Matthew C. MacArthur, a Ph.D. candidate at Stanford University, and extend and enhance a model that was already implemented and had incorporated many of the structural features that we were interested in.<sup>16</sup> MacArthur's original model attempted to explore the responses of a set of political actors (voters) to various government policies. The political actors either share the government's policy position or they do not. Those that do not eventually express their disagreement by attacking the police and/or those citizens who endorse or at least accept the government's policies. The government may tolerate some dissent, but the more violent forms of dissent are punished, if possible, by the police. The adoption, adaptation, extension, and enhancement of MacArthur's Java-based model turned out to constitute the second major stage of our technical approach, namely to develop an executable agent-based model and simulation.<sup>17</sup>

---

<sup>13</sup> <http://www.brook.edu/dybdocroot/es/dynamics/models/>

<sup>14</sup> <http://www.mcw1.usmc.mil/divisions/albert/index.asp>

<sup>15</sup> (Lauren 2002)

<sup>16</sup> (MacArthur 2002) Mr. MacArthur graciously allowed us to use and extend his Java-based model and has been working closely with us to extend the model in directions that are mutually useful.

<sup>17</sup> Some justification for the use of Mr. MacArthur's model as the foundation for our work may be needed. Selection of MacArthur's model (to be called M) for our preliminary work was based on several factors, the most important of which is that it appeared, on the face of it, to capture one fundamental aspect of terrorist organization behavior: terrorism is a political phenomenon; it is the use of terror to (attempt to) bring about political change. In general, the M model of dissent and rebellion attempts to capture and reflect this fundamental property of terrorist behavior. Another important reason for adopting (and then adapting) the M model is that it includes the key notion of alliances: dissenting individuals can ally with other dissidents, forming an alliance whose collective power can be appealed to by a member of the alliance in the individual's expression of dissent (generally attacks on other non-alliance individuals). It is easy to see how these "alliances" can be thought of as "terrorist networks" and how the conditions for forming such alliances might be extended to include the need for various potential network members to have certain capabilities (e.g., financial, leadership, communications, etc.). Another important feature of the M model is its potential for a more full-blown form of adaptation with individuals or alliances self-modifying their rules of behavior based on global objective functions. If the overall objective of the terrorist networks are not being achieved, the organization as a whole can change its behavior, deciding to attack different targets, perhaps, or altering its basic constitution by changing the criteria for membership in the group, etc. Finally, we decided to use the M

The third stage of the technical plan was to systematically run the model to generate simulation results that could be studied and analyzed in an attempt to answer the questions posed for this research effort. Given that we now had an agent-based model of a complex adaptive system (interpreted as a terrorist organization), could anything of value be learned from the simulations produced with the model? And, more generally, did this ABM approach offer advantages to an anti-terrorism<sup>18</sup> analyst, let's say, that other, more traditional techniques could not readily provide?

The fourth and final stage of the project was to report on our research findings and conclusions. This paper is the culmination of this final phase of the project. Unfortunately, it can only be considered a preliminary report, given the need for additional work in extending and enhancing the model, and the need to generate and analyze more simulations.

## 1.5 Future Applications

Although the specific asymmetric threat domain we decided to focus on was terrorism, it is not hard to imagine how the approach and techniques can be applied to other issues, including local insurgencies, illicit drug trafficking, foreign counter-intelligence, law enforcement—specifically organized crime, street gangs, radical militias—, and ethnic violence. These concerns share many of the characteristics of terrorist organizations that make the latter so difficult to contain, including: complex adaptive behavior, illicitness or illegality—and hence covertness, violence, and asymmetry with respect to means. One might also argue that ABM and similar techniques may be more appropriate to one or

---

model because it was available, executable, and reasonably robust. It works, but could be modified easily and extended to suit our particular interests.

<sup>18</sup> Technically, one can distinguish between anti-terrorism and counter-terrorism. “Anti-terrorism activities are those activities that happen before a terrorist event occurs and are aimed at preventing the incident. These are primarily law enforcement, intelligence, and investigative types of measures. They could include hardening of targets, training of personnel, public awareness and other security measures. Counter-terrorism activities are those activities that happen after a terrorist attack occurs. This involves the response to, the mitigation of, and the recovery from a terrorist incident” (<http://www.msfc.org/Legislative%20Page/Position%20Paper.htm>). Since the focus of this research is on the prevention side, we should refer to anti-terrorist activities rather than counter-terrorist actions, but, since “counter-terrorism” is well entrenched in the public’s argot, we’ll use that term throughout the paper.

more of these other domains. The approach might be more useful in addressing the problem of street gangs, for instance, given their geographical specificity.

## **1.6 Definitions**

Definitions of some key concepts are in order. Terrorism is the specific asymmetric national security threat which we used as the target domain in this work. Complex adaptive systems are those that combine systems behavior with complexity science and adaptation. Agent-based modeling (and simulation) is the particular information technology that seems especially suited to the analysis of complex systems, whether adaptive or not. Social network analysis is another information technology technique being used to explore the systems dynamics of organizations and social networks. These terms are explained more fully below.

### **1.6.1 Terrorism**

There are literally hundreds of different definitions of terrorism. The Federal Bureau of Investigation (FBI) defines it as "[t]he unlawful use of force against persons or property to intimidate or coerce a government, the civilian population, or any segment thereof, in furtherance of political or social objectives."<sup>19</sup> Christopher Harmon defines it similarly as "the deliberate and systematic murder, maiming, and menacing of the innocent to inspire fear for political ends".<sup>20</sup> It is "[n]ow a well-established feature of world politics and conflict, [and] is used by single-minded small groups, state agents, and broader insurgent movements to seek political and military results judged difficult or impossible to achieve in the usual political forums or on the battlefield against an army. Terrorism is always political, even when it also evinces other motives, such as the religious, the economic, or the social."<sup>21</sup> "But while all terrorism has a political purpose, it certainly is distinguish-

---

<sup>19</sup> <http://www.msfca.org/Legislative%20Page/Position%20Paper.htm>, November 2, 2003.

<sup>20</sup> (Harmon 2000, p. 1) The original definition was adopted by The Johnathan Institute in Jerusalem at a 1979 conference.

<sup>21</sup> (Harmon 2000, p. 1)



able—technically and morally—from civil dissidence, other forms of civil violence, or revolution, which are also political phenomena. All the others are possible...without terrorism.”<sup>22</sup>

Several themes are woven into these characterizations of modern terrorism. The modern terrorism organization is often decentralized, transnational or state-less, global in its reach, and asymmetric. They are asymmetric in the sense that they are relatively small and militarily weak. They cannot confront the governments they oppose directly so they resort to other, indirect means, specifically terrorism. Osama bin Laden can’t defeat the US directly and can’t drive the US out of the middle-East with military, diplomatic, or political means, so he relies on suicide bombers. Terrorism is an asymmetric threat employed by political entities that lack the political, economic, or military means to confront its adversaries directly or head-on. Instead, they resort to means that can be leveraged in an attempt to effect their goals. Terrorism is a low-cost but high-impact means of attack. A series of suicide bombings can inflict an emotional toll that exceeds the financial cost of executing a series of well-planned attacks against civilian populations. As Wu Ch’i is reported to have said: “One man willing to throw his life away is enough to terrorize thousands.”<sup>23</sup>

Al-Qaeda is global in its reach in that it has mounted successful attacks in Africa, North America, Europe, the middle and far-East, and Central Asia. It and other modern terrorist organizations are transnational or state-less, meaning that they are not dependent on financial or other support (e.g., territorial sanctuary or training sites) from sovereign nation states. Finally, they are decentralized in the sense that there does not appear to be a rigid and disciplined centralized and hierarchical command structure. Rather they are composed of “cells” consisting of a handful of operatives that act somewhat autonomously and independently of any central leadership.

---

<sup>22</sup> (Harmon 2000, p. 1)

<sup>23</sup> <http://faculty.ncwc.edu/toconnor/429/429lect01.htm>

It's important to note that while terrorism seems often tactical in nature, it is fundamentally a strategy that is "the considered application of means to advance certain ends."<sup>24</sup> "[T]errorism is chosen for definite purposes; it is a chief means to advance political ends. While sometimes it is the *only* well-developed means a given group employs, successful terrorists invariably use additional means and forms of effort: political, social, military, or even humanitarian."<sup>25</sup>

"Because terrorism is a strategy, it may be used by different groups or governments for very different ends. It often serves multiple purposes: the same kidnapping can be intended to shock the public, to cripple a politician, and to raise operating funds, all at once."<sup>26</sup>

"Special characteristics set terrorism apart and bear upon its use as a strategy. Unlike political parties and their fights, terrorist activities are illegal, and explicitly anti-legal. Yet, unlike typical crime, or transnational organized crime, terrorism is directed at a public purpose; its money making is incidental to that end.... Terrorists distinguish themselves by targeting civilians, the unarmed, and the innocent.... Terrorists aim to shred the status quo and make a 'new order' of their own design. Most use of force in non-terroristic settings aims to destroy or reduce an identified threat; the very essence of terrorism is in the calculated use of violence to spread alarm through a wider audience; the actual target may be almost incidental to the desired effects, which expand outward like shock waves."<sup>27</sup>

"The charters and communiqués that the [terrorist] groups write, and the interviews they eagerly arrange with friendly newspapers, are as instructive as the incendiaries they leave

---

<sup>24</sup> (Harmon 2000, p. 44)

<sup>25</sup> (Harmon 2000, p. 44)

<sup>26</sup> (Harmon 2000, p. 44)

<sup>27</sup> (Harmon 2000, p. 44f)

in department stores or under diplomats' cars.... Most often, terrorism is not mindless; it is the calibration of violence and fear for political effect.”<sup>28</sup>

“Typically, terror groups envision themselves as locked in an unlimited effort towards an extraordinary end, such as total anarchy, or the revolutionary seizure of full political power, or the creation of a holy and perfect religious kingdom on earth.... The usual extremism of the ends bears directly upon choice of means.... [Strategists] view related strategic questions—such as whether to halt attacks, or mix in legitimate politics, or ally with organizations in nonviolent struggles—as prudential issues. They do them all if they calculate that they will be useful, just as they carry on the ‘deliberate, systematic murder, maiming, and menacing of the innocent to inspire fear and gain a political end.’”<sup>29</sup>

It should be noted that the characterization of terrorism provided above is not universally accepted. Sun-Ki Chai has argued that terrorism is basically an intra-organization technique for “antigovernment leaders” [read “terrorist group leaders”] to “solve a number of [intra-group] informational problems: imperfect monitoring of cells, difficulty in communicating examples of punishment, and difficulty in communicating the organization’s activities. Violence solves these problems in ways that do not threaten the secret and disperse structure necessary to maintain an underground antigovernment organization.”<sup>30</sup> Chai’s view is an attempt to reconcile the “conventional” theory—terrorism is aimed at effecting political change—with the fact that of “335 ‘episodes and campaigns’ from 1961 to 1977 involving ‘terrorism’ [defined as] politically motivated violence conducted by clandestine groups...only a handful...appear to have had any lasting effect on national political systems.”<sup>31</sup> In other words, if terrorism is a strategy for effecting political change, how do we account for its remarkable failure? Implicit is the question, If terrorism is so ineffective, why does it still prevail as the strategy of choice for so many anti-

---

<sup>28</sup> (Harmon 2000, p. xviii)

<sup>29</sup> Harmon 2000, p. 73

<sup>30</sup> Chai 1993, p. 108

<sup>31</sup> Chai 1993, p. 100

government groups? The implicit answer, of course, is that terrorism may not be a political weapon at all.

### 1.6.2 Complex Adaptive Systems

A complex adaptive system (CAS) is a type of system—albeit a complex one—capable of adaptation. A *system*, is “a set of units or elements...interconnected so that changes in some elements or their relations produce changes in other parts of the system.”<sup>32</sup> A system is *complex* if its interacting elements interact in a non-simple way with “the entire system exhibit[ing] properties and behaviors that are different from those of the parts.”<sup>33</sup> An *adaptable* system is one whose behavioral repertoire changes in order to remain viable in an environment. (Note that a system does not need to be complex in order to adapt to its environment. Simple single-celled, asexual prokaryotes have been found in every environment on the earth and account for most of its biomass.)

John Holland’s *Hidden Order: How Adaptation Builds Complexity* (1993) is the *modern* classic treatment of complex adaptive systems, but the notion of a CAS is hardly new. Indeed, one needs to go back to the early 1940s to find the beginnings of a vast literature that has come to be known as “modern systems theory” (Buckley 1998, p. 12). The predominant work in the analysis of CASs using ABMs, however, is relatively recent. Swarm, for instance, an early platform for ABM development, was begun in 1994 by Chris Langton at the Santa Fe Institute in New Mexico. This is no doubt due not only to the growing popularity of object-oriented programming languages such SmallTalk, C++, and Java, but also the immense computational power becoming available at relatively low cost to the desktop.

---

<sup>32</sup> (Jervis 1997, p. 6)

<sup>33</sup> Ibid. Compare this definition to a similar one by Herbert Simon: “Roughly, by a complex system I mean one made up of a large number of parts that interact in a nonsimple way. In such systems the whole is more than the sum of the parts, not in an ultimate, metaphysical sense but in the important pragmatic sense that, given the properties of the parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole” (Simon 1981). For more on complexity, see below, Section 1.6.4.

### 1.6.3 Agent-Based Modeling

Agent-based modeling (ABM), in the sense in which we are interested here, is the attempt to model physical world phenomena with software agents. Software agents are software representations of physical world entities that, in the physical world, would be relatively autonomous and heterogeneous. They are autonomous in that they act in some sense of their own “volition,” following internal rules of behavior based on internal motivations and beliefs about the world. These physical world entities, and their *in silico* counterparts, are heterogeneous in that they differ in the rules they follow, in their motives, and in their beliefs about the world, either the physical world for physical world entities or the synthetic world that comprises the environment of an agent-based model and simulation. Note that this characterization of ABM refers to *representations* of physical world entities. ABMs are supposed to be models of the physical world and are supposed to provide some physical world benefit, either a better *understanding* of natural phenomena or, ideally, a way to *predict* physical world events.

How might ABM be used to counter asymmetric threats such as transnational terrorism? As sketched above, the core idea is to develop an ABM that can model and simulate various facets of terrorist network behavior. An examination of the simulation’s behavior may afford better insight into actual terrorist behavior, perhaps exposing weaknesses in the simulated network that could be exploited in the physical world. It is even possible that the simulation could point to future events that could even be prevented. The use of ABM techniques—and, in indeed, modeling, in general—represents an argument from analogy. Model  $M$  is analogous (in all relevant respects) to physical world behavior  $W$ ;  $M$  exhibits (simulated) behavior  $m$ ; therefore,  $W$  (must) exhibit (physical world) behavior  $w$ , where  $w$  is the physical world analogue to  $m$  in the simulation. So, for example, if an ABM simulation of terrorist network behavior indicates that terrorist attacks only occur under certain conditions, then one can infer, using the argument from analogy, that physical world terrorist attacks can occur only under certain conditions. The hope is that the ABM will tell us precisely what those conditions are.

An ABM is seemingly an ideal way to investigate complexity. Complexity is the macro-level behavior exhibited by various systems generally composed of hundreds, thousands,

or even millions of individual parts, each of which are determined by usually a small number of rather simple rules or laws.<sup>34</sup> Complexity is conceptually orthogonal to both uniform (i.e., simple, repetitive, or nested) behavior on the one hand and completely random behavior on the other.<sup>35</sup> It is behavior that is not random nor is it simple and easily repeatable (and predictable). Complex behavior is inherently—almost by definition—*non-predictable*. This non-predictableness of complex systems, in general, and ABMs (and simulations) designed to emulate physical world complex systems, in particular, presents the modeler, however, with a problem. Exactly how can ABMs be of any value (in exploring complex systems) if complex systems are inherently non-predictable? If the complexity-oriented ABMs cannot be used to predict model behavior, how can they possibly predict physical world behavior? And if they can't be used to predict physical world behavior, how is it that they can offer any value to those who ultimately have to take the battle into the physical world? Put simply, if terrorist networks are complex adaptive systems, then they are inherently—by definition—non-predictable in their complexity-based behavior. Modeling and simulation, in general, and agent-based models and simulations, in particular, are useful in proportion to the extent they allow us to predict physical world behavior based on model and simulation behavior. But if the ABM actually captures complex behavior, it too is inherently non-predictable and, consequently, of no value in predicting physical world behavior.

There are a few ways of countering this argument. One line of response often heard is that ABMs are intended to provide *insight* into the behavior of the complex systems they are intended to model. Exactly what this discovered insight is, however, is more often than not hard to discern.

---

<sup>34</sup> This is only a rough and ready definition of “complexity.” Seth Lloyd, Professor of Mechanical Engineering at MIT, has collected a list of almost 40 “measures of complexity” (<http://web.mit.edu/esd.83/www/notebook/Complexity.PDF>).

<sup>35</sup> Stephen Wolfram distinguishes four classes of (non-random) behavior of cellular automata arising from random initial conditions: very simple (class 1), where “almost all initial conditions lead to exactly the same uniform final state;” simple structures that may or not repeat (class 2); seemingly more complicated, and, in some respects, almost random but with “small-scale structures” seen “at some level” (class 3); complex, involving “a mixture of order and randomness” (class 4) (Wolfram 2002).

A less fuzzy response is to argue that ABMs are not intended to be predictive. There are always too many unknowns, too many variables to enable a model/simulation to predict anything—even in the non-complex world—but a range of probable outcomes, given a wide range of values for the input parameters, is possible and of value.

The problem with this argument is that it tries to have it both ways; it tries to say that such models are not intended to be predictive, but then turns around and says that they are predictive—within a probability range. “I can’t say exactly what will happen,” the proponent argues, “but there’s a 90% probability that this will happen.” But this is still a prediction and perhaps a valid one and the kind made everyday by one’s local weather forecaster (“there’s a 90% chance of rain tomorrow”). There’s nothing wrong with this except the argument is beside the point if it is attempting to defend the value of ABM tools for the analysis of complex systems. Such systems are inherently non-predictable. Left to run a little longer, the simulation (used to forecast the 90% chance of rain) might easily undergo a “phase-change,” utterly destroying the basis of the original prediction. All that can be claimed legitimately for an ABM that genuinely models (and is used to simulate) complex system behavior is that *if* one has succeeded in modeling the important elements (i.e., the actors and their behavior) of the physical world system and *if* these initial model/simulation conditions ever occur in the physical world, and *if* one can register physical world time periods with each iteration of the simulation, *then* this is the behavior that will occur—that can be validly predicted to occur—after a physical world period of time corresponding to so many iterations of the simulation. This rationale is not all unreasonable. (Given the number and nature of the qualifications, no wonder.) But the salient question then becomes, To what extent can these qualifications ever be met? Are there any physical world situations in which all and only the important initial conditions can be known, in which all and only the important actors can be isolated and wholly described (in all relevant, important respects), and in which the temporal flow can be sufficiently defined to enable temporal registration (or calibration) of the simulation? And, if so, is this physical world situation of any real interest?

Instead, we would like to suggest another approach to the justification of ABMs and simulations of complex behavior that avoids the problem of inherent non-predictability of

complex systems. Simply stated, an ABM (and simulation) of complex<sup>36</sup> adaptive systems may reveal novel and surprising forms of behavior that could be of tremendous value when looked at in light of the physical world behavior or events. An ABM/simulation of the al Qaeda network, for example, that had operatives enrolling in flight schools could have alerted authorities to look for similar physical world events even though the ultimate purpose of such actions were not clear. It is this, we suggest, as both the immediate and long-term value of ABMs/simulations of *adaptive* systems, whether known to be complex—in a formal, rigorously defined sense—or not. The emphasis is on the emergence of *novel* behavior in the simulation, on behavior that is on the face of it unexpected, and perhaps explicable only by an examination of the behavior that led to or engendered the novel behavior. The hope (goal) is that such examinations would reveal important system-level principles that came to (partially) govern the simulation at that time (and which might quickly go away, being replaced by other “emergent” principles).

We are suggesting that the possible value of ABM is to presage potential terrorist behavior that is both (1) novel (and, most likely, unpredictable) and (2) amenable to detection in the physical world. We are suggesting that ABM technology could be used to “explore” or “search” a “space” of “potential behavior” of terrorist networks. We are looking for “behavior” that is “novel” or “unexpected.” That is to say, we are looking for behavior in our simulated terrorist networks that is beyond the readily understood, beyond the ken of the anti-terrorist community. It is not clear how this will be accomplished, but the basic idea is simple enough. We would like to have an ABM simulation in which behavior emerges that is “unexpected” and, initially, at least, inexplicable. Such behavior—if it poses a potential threat—will prompt further analysis and ideally occasion the search for suitable indicators of similar physical world behavior. This particular desideratum—basically an “actionable” result—serves as a constraint on the overall model.

---

<sup>36</sup> In some sense, the phrase “modeling a complex system” begs the question. Complex systems need to be discovered. One can assume that certain systems (e.g., global weather) are complex systems. Other systems which we might suspect are complex may turn out not to be so. Determining if a



#### 1.6.4 Complexity Theory

The above characterization of ABM suggests how the technology may be employed to address a rather specific problem. The characterization does not underscore why ABM has been singled out for consideration rather than other modeling—or indeed other—analytical techniques. What then is so special about ABM?

ABM is a branch of *complexity* theory. Complexity theory is based on the assumption that (natural) phenomena of any real interest are inherently *non-linear*, that is, characterized by extreme sensitivity to initial conditions, subject to abrupt changes in macro-level behavior, and inherently non-predictable. Complex or non-linear behavior stands apart from regular, repeatable behavior on the one hand and completely random behavior on the other. Complexity theorists believe that much of the behavior observed in the world—the weather is a frequently used example—is certainly not regular, but it is not random either: it is complex and thus subject to whatever constraints the natural, physical laws of complexity imposes upon it and which the complexity theorists are trying to discover. But perhaps the most important impetus for using ABM technology to study complex adaptive systems is that it's possible to do so due to the advances in performance and versatility afforded by modern information technology. It is not only possible but relatively easy to run complex simulations thousands of times with different parameter values in order to get a better understanding of the combined effects of perhaps literally billions of interactions within the model. While this has always been possible since the advent of the digital computer, it has only been within the last decade or so that it has become practical and affordable on the desktop with a single-user machine. The modern digital computer has given us new ways to methodically and systematically explore very large possibility spaces. The ways in which these explorations are carried out, however, have to be justifiable, and the interpretations of any discoveries in terms of their applicability to the physical world have to be sound. We'll come back to each of these issues later, in Section 4.

---

model/simulation is of a genuinely complex system is not so easy. See Chapter 10 of (Wolfram 2002), and (Pigliucci 2000).

Sufficiently rich ABM simulations seem to exhibit similar complex behavior. From many interacting individual software agents, each following rather simple rules of behavior, non-predictable “patterns” often emerge. This *emergent pattern* phenomenon is almost the *sine qua non* of complex systems and ABM simulations. Unexpected regularity—which isn’t really regular nor predictable—emerges from the complex interplay of simple forces or individuals. If ABM simulations exhibit this surprising kind of behavior, they must be the proper tool for investigation complex behavior in the physical world.

Complexity theory is the attempt to organize and guide the study of complex interactions and the emergent<sup>37</sup> properties they engender. Complexity theory subsumes chaos theory (i.e., all chaotic systems are complex, but not vice versa). Complex systems are deterministic, but not predictable (or very difficult to predict). Complex behavior (or chaotic dynamics) are usually the property of non-linear systems, that is, systems described by sets of non-linear equations, but not always. If a system is chaotic, it is non-linear, but the converse does not hold. If a systems is non-linear, it is not necessarily chaotic. More precisely, if one considers  $N$  nodes with  $K$  possible connections, if the ratio of  $N$  to  $K$  is large (many nodes but few connections), the system is called sub-critical and is amenable to classical analytical techniques in mathematics, physics, and biology. If the ratio of  $N$  to  $K$  is approximately even, the system is critical (the “edge of chaos”). If  $K$  exceeds the number of nodes,  $N$ , the system is called super-critical and is chaotic. Complexity theory contends itself with the “critical” system and the transitions between sub-critical to critical and from critical to super-critical.<sup>38</sup>

---

<sup>37</sup> Thw notion of “emergent” phenomena needs to be scrutinized. The term is used as if it’s well understood. We read recently vis-à-vis “boid” (cf. Reynolds 1987) behavior that “flock obstacle avoidance” is an emergent behavior, not predictable from knowledge of the simple rules that govern the behavior of the individuals in the (simulated) flock. Since there were no rules governing flock behavior, the fact the flocks of boids avoided obstacles was deemed “emergent” and, given the enthusiasm of those calling our attention to such behavior, practically miraculous and worthy of awe. It’s almost as if this “miraculous” (“mysterious”) emergent behavior is in itself enough to justify the expenditure of large sums of DoD R&D money to apply this interesting behavior to the real world.


<sup>38</sup> (Pigliucci 2000)

Remember the original question: What is afforded by analyzing asymmetric threats as complex adaptive systems? We can safely assume that terrorism is an asymmetric threat. Terrorism is an asymmetric *means* to achieve some goal. Terrorists don't resort to terrorism because no other means are available. Adversaries resort to terrorism because they are in an asymmetric position vis-à-vis their adversary, and "terrorism" simply becomes the term used to describe that behavior. But what of the assumption that such adversaries are complex adaptive systems? Again, it's probably safe to assume that they are systems (a collection of interacting parts), and they are most likely adaptive as well. The "system" is "intelligent" enough to change its behavior in response to the behavior of its adversaries. The important question that remains then is whether the system is "complex" or not, and if so, whether ABM is the best way (or even a good way) to analyze its behavior with the intent to defeat it? Are terrorist organizations complex systems? Is ABM a legitimate technique—whatever else its merits may be—for analyzing complex systems and their behavior?

Let's address the question of the "complex" nature of terrorist organizations. What exactly does this question mean? Are we to think of a terrorist organization as a complex system whose "internal" behavior gives rise to emergent phenomena? And is this emergent phenomena necessarily expressed as terrorist attacks—acts of terrorism? Or might the emergent phenomena be expressed as something more benign, perhaps the dissolution of the group itself? Perhaps in saying that a terrorist organization is complex means that the time-series plot of the group's terrorist acts are "chaotic," that is, deterministic but not predictable. (Or, more generally, that all acts of terrorism are inherently unpredictable given that such acts are the expression of chaotic—and, by definition, non-predictable—behavior.) This raises a quandary. If terrorism is complex (and/or chaotic), then it is inherently non-predictable (its underlying deterministic nature notwithstanding). And if non-predictable, then why bother? We can't do anything anyway. But if terrorist organizations are not complex systems, then what is the value of using ABM techniques—techniques well suited for—the analysis of complex systems? ABM may be of immense value in analyzing the *complicated* behavior of non-simple systems. It is not obvious that such techniques are of value in analyzing the behavior of *complex* systems as they give rise to *emergent* phenomena. But to return to the question raised above, is terrorist or-


ganization behavior genuinely complex (chaotic)? And how would we determine it? One technique is to look at a phase plot of data generated by the hypothesized complex system. Whereas a simple time series depiction of a systems behavior will look random, its phase plot can sometimes reveal complex system behavior

We might plot a time series of terrorist events as in Figure 1.



**Figure 1. Total International Terrorist Attacks, 1981 – 2002<sup>39</sup>**

Figure 2 shows a detrended version of this data, that is, a graph of the variation in the number of incidents rather than the absolute numbers.




**Figure 2. Detrended Data**

Finally, Figure 3 shows a phase plot of the detrended data from Figure 2.

---

<sup>39</sup> From Appendix H, Statistical Review, Patterns of Global Terrorism, U.S. Department of State.



**Figure 3. Phase Plot of the Detrended Data**

**Figure 3** does not show any obvious patterns that might suggest chaotic as opposed to random behavior. This may be due simply to the paucity of data points (21).

### 1.6.5 Social Network Analysis

According to Valdis Krebs, a leading researcher in the field, “[s]ocial network analysis (SNA) is the mapping and measuring of relationships and flows between people, groups, organizations, computers or other information/knowledge processing entities.”<sup>40</sup> Using traditional link-node (or edge-vertex) graphs, where the “nodes [vertices] in the network are the people and groups while the links [edges] show relationships or flows between the nodes,” the social network analyst can graphically depict and mathematically analyze the structure of complex social networks. Some of the key concepts used in SNA include degrees, betweenness, closeness, and centrality.<sup>41</sup>

*Degree* is a measure of network activity in terms of the number of direct connections a node has. Generally, the more connections, the better. But not always. Krebs writes that “[w]hat really matters is where those connections lead to—and how they connect the otherwise unconnected!” *Betweenness* refers to the often crucial role a node may play as po-

---

<sup>40</sup> (Krebs 2002)

<sup>41</sup> Ibid.

tential broker in a network because, while limited in the number of direct connections, it serves to connect important network constituencies. “A node with high betweenness has great influence over what flows in the network.”<sup>42</sup> *Closeness* refers to the number of direct and indirect links a node has to all other nodes in the network. Nodes with a smaller number of such links connecting them to all other nodes of the network are closer than those nodes with a larger number of links. Such nodes “are in an excellent position to monitor the information flow in the network—they have the best visibility into what is happening in the network.”<sup>43</sup> Social network *centrality* refers to the degree to which a network “dominated” by a “few well connected hubs [that] can fail abruptly, that is, split into unconnected sub-networks—if those nodes are disabled or removed.”<sup>44</sup> Networks with low centrality—those not dominated by one or a few nodes—are “resilient in the face of many intentional attacks or random failures.”<sup>45</sup>

Sun-Ki Chai’s theory of antigovernment violence based on organizational economics—that is, on the theory of “how individual incentives and processes within organizations affect collective outcomes”<sup>46</sup>—gives credence to the value of using SNA tools in the analysis of terrorism.

Some questions that still need to be addressed include: Is anyone using ABM to really explore the *validity* of SNA? How sound are the concepts of “centrality,” “betweenness,” “cognitive centrality,” etc.? How are they applicable to anti-terrorism? Could we use ABM techniques to explore how terrorist organizations might naturally evolve to a particular kind of “social network” as they seek to achieve their goals and avoid disruption/dissolution by adopting various network structures that emphasize (or de-emphasize) these SNA concepts? Is a terrorist network that avoids to the maximum extent possible

---

<sup>42</sup> Ibid.

<sup>43</sup> Ibid.

<sup>44</sup> Ibid.

<sup>45</sup> Ibid.

<sup>46</sup> (Chai 1993)

nodes exhibiting high-centrality a network that is particularly resistant to typical anti-terrorist attacks?

## **1.7 Organization of the Report**

The remainder of this report is organized in two main parts. The first part, Sections 2 and 3, describes the model and simulations. The second part, Sections 4, addresses the overall objective of this research, namely, what is the potential value of ABM for the analysis of CAS, in general.





## **2. Terrorist Organizations Views Complex Adaptive Systems**

### **2.1 Basic Idea of the Model**

The model we developed for this project is intended to generate “terrorist organization-like” behavior. Certain model entities represent terrorists. Other entities stand for ordinary citizens. Still other model actors play the role of the (anti- or counter-terrorist) authorities. Each of these three types of actors behaves according to various rules of behavior. Some of these rules are shared by all the actors or agents in the model. For example, the rules governing agent movement are the same for all actors. Other rules of behavior are unique to the different kinds of agents. Terrorists mount attacks on innocent citizens and the authorities, something that normal citizens never do. The rules governing actor behavior may change during the course of a simulation, usually in an attempt by an actor group to adapt to a changing environment.

The basic idea behind the model is that if the set of rules chosen to direct agent behavior are approximately correct at the individual actor level, then aggregate-level behavior should reflect the behavior we encounter in the real (i.e., non-simulated) world. If not, then our choice of behavior rules is probably mistaken.<sup>47</sup> This may seem obvious, but this fact highlights an important aspect of this type of modeling and simulation effort.

It would not be unfair to say that we neither know what we’re trying to model/simulate nor how to do it—except in a broad sense. We don’t presume to know exactly what motivates a terrorist or exactly how terrorist groups are organized and operate. It’s likely that no one really knows, not even the terrorists themselves. Nor, of course, are motivations and ways of organizing and operating invariant between terrorists or terrorist groups. We

---

<sup>47</sup> Not necessarily, of course. There are many other factors that may have nothing to do with the rules governing individual agent behavior that affects the aggregate-level behavior of a simulation. The order in which the behavior rules are invoked for each actor may affect the outcome. The procedure for selecting each agent for behavior execution may affect the simulation in significant ways. These endogenous “artifact” effects can be detected, however, and techniques (e.g., randomization of control flow) can be used to mitigate—if not eliminate—their influence.

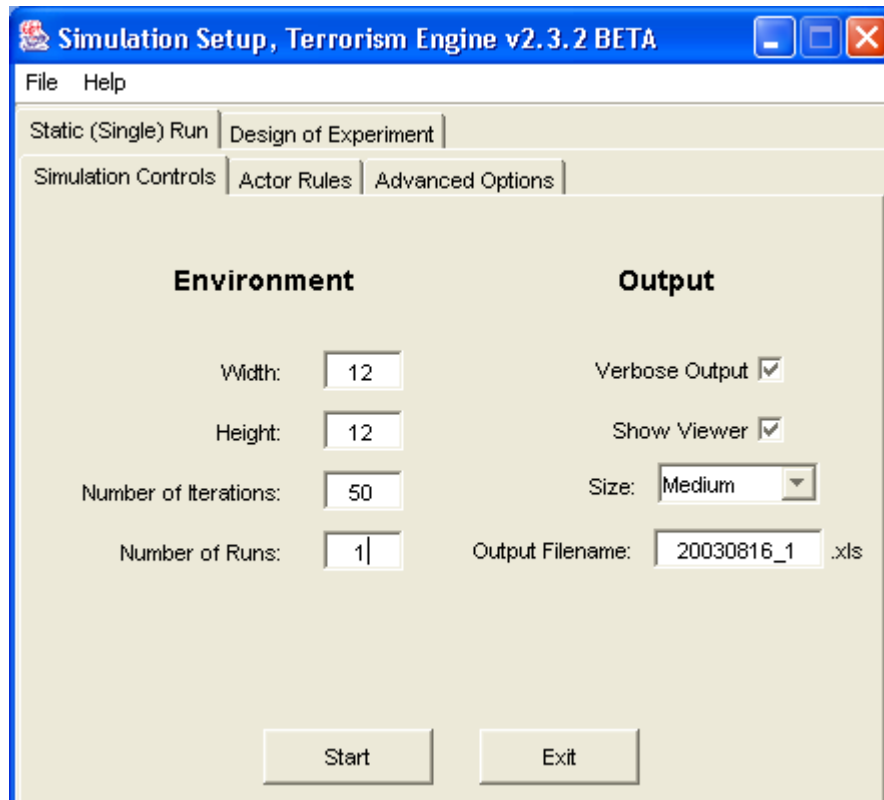
do have at hand, however, various theories of terrorist behavior, and one way we see to test these theories is to use agent-based modeling and simulation techniques. If a theory of terrorism is sufficiently precise to allow simulation in a computer-based model, then it should be possible in principle to verify the plausibility of the theory by conducting simulations that can be compared to physical world events. To the extent that a simulation accords with actual events, the original theory is (partially) verified: the simulation experiment provides some evidence in support of the theory. Unfortunately, even this hopeful line of investigation is beyond the reach of the present work. There are no theories of terrorism—or, indeed, of any complex human social behavior that we are aware of—that are sufficiently formalized to permit simulation testing. This point will be returned to later.

## **2.2 Basic Model Parameters**

The easiest way to describe the basic parameters used in the model is to work through the options available in the graphical user interface, the initial screen of which is depicted in Figure 4. There are two basic modes in which simulations may be conducted, a “Static Mode” and what we call the “Design of Experiment” mode. Under the first mode, a user provides a static set of model parameters to be used during a simulation, either a single run or multiple runs. Variations in the behavior of the model in the static mode is due entirely to the stochasticism introduced by the use of a (pseudo) randomizer for actor placement on the landscape and other places where a randomizer function is obviously called for. Under the Design of Experiment mode, model parameters are automatically and systematically changed in accordance with user specification for each separate simulation. This simulation mode is designed to enable the analyst to systematically “sweep” the parameter space to look for those parameters that have the most effect on the overall behavior of the model. It is this mode that enabled us to produce the results presented in Section 3 below.

With the Static Mode tab of the Graphic User Interface (GUI) active, three sets of parameters are selectable. “Simulation Controls” are used to specify the size of the environment (the landscape) as well as mechanical controls that are materially irrelevant to the behavior of the model as a model of complex adaptive system behavior. The “Actor

Rules” tab allows the user to specify the initial number of actors by actor type as well as the various thresholds to be used in the actor transformation rules. Within the “Advanced Options” tab, the user can set “sympathy” thresholds, whether initial values are determined by a uniform or Gaussian distribution function, whether a seed should be used for the random number routines to enable result repeatability, and whether “learning” (or “adaptation”) is to be enabled or not.



**Figure 4. Initial Window of the Graphical User Interface**

### **2.2.1 Simulation Controls**

There are two sets of simulation controls. One set controls the basic environment. The other controls display aspects and the capture of simulation output.

### 2.2.1.1 Environment

#### **Width and Height**

The agents of the model move about and interact with each other on a rectangular “landscape,” the dimensions of which can be specified by the user. In Figure 4, the landscape dimensions have been set to use a square matrix (or lattice) of 12 units to each side. The landscape size sets an upper limit to the number of possible actors. The user determines the size of the landscape by entering values for width and height.

For purposes of agent movement and interaction, the landscape is considered a torus such that each square has by definition eight neighbors, four orthogonal and four diagonal. The upper-most left-hand square of the matrix has the eight neighbors indicated in the diagram of Figure 5, based on a 12-unit by 12-unit matrix.

12,12	1,12	2,12
12,1	<b>1,1</b>	2,1
12,2	1,2	2,2

**Figure 5. Landscape as Torus**

#### **Number of Iterations**

The user may specify the number of iterations of the central loop of the model for each simulation run. Although in Figure 4 the number of iterations is set to 50, a minimum of 200 iterations per run is generally used. An iteration is a single step in the simulation during which all actions that are available to execute at that step are executed. In this model, every existing actor is “processed” during each iteration, that is to say, all actor behavior rules are applied to each existing actor during each iteration. The order in which actors are “processed” during each interaction is determined by their position on the landscape. The fundamental execution loop in the model simply looks at each landscape element in a regular, systematic way: left-to-right, top-to-bottom.

### Number of Runs

The user may specify the number of distinct runs to be executed during one simulation. In Figure 4, the number of runs is set to 1. In this particular case, the single run will consist of 50 iterations. Using multiple runs and taking the average (or some other statistically relevant) value of salient aspects of the simulation strengthens the claim to statistical validity of the results.

#### 2.2.1.2 Output

There are four controls available for adjusting output parameters during a simulation.

##### Verbose Output

An execution trace of certain key events is always written to a file during each simulation run. This trace file (named “commandLineOutput”) always captures basic information such as date and time of the run and initial parameter values. Figure 6 is a portion of a typical output file.

A more verbose form of the execution trace can be displayed to the command line window. This control can be toggled on-or off by checking the verbose output check box.


```
RUN: 1
OpenCyc Error java.net.ConnectException: Connection refused: connect
Aborting Run.
ITERATION: 1
Terrorist moves to find allies.
Competence = 0.0
Since last change = 1
Worst = 9.999999999E8, Current = 0.0
Competence = 0.0
Since last change = 1
Worst = 9.999999999E8, Current = 1.0
```

**Figure 6. Output File Example**

##### Show Viewer

The show viewer control is used to toggle on-or-off a graphical display of the simulation. It is used mostly for debugging and to display the basic behavior of the underlying rules. It is essentially irrelevant to the model and may be suppressed for performance purposes. A typical viewer display is depicted in Figure 7. Each colored cell represents an agent:

green for passive citizens, blue for police, orange for actual (but unknown to the police) terrorists, red for known terrorists, and yellow cells for latent terrorists. The black lines indicate various terrorist networks. One network consists of the three squares connected to the agent in the upper right-most corner.



**Figure 7. Viewer Display**

### **Size**

The output size menu allows the user to adjust the viewer display to small, medium, or large. It is non-operative if the show viewer is un-checked.

### **Output Filename**

The output filename is the name of the output Microsoft® Excel spreadsheet file used to capture the results of a simulation run (or runs). The file name encodes both the date on which the file was created and the particular ordinal instance of the run on that date. The

file captures the precise date and time of the simulation run, the parameter settings in effect upon initialization, and a number of important simulation values at each iteration. These values include:

- Number of known terrorists (i.e., terrorists known to the police)
- Number of terrorists (i.e., number of all terrorists, known and unknown)
- Number of latent terrorists (i.e., number of terrorists who have not yet entered into a terrorist network)
- Number of police
- Number of events (i.e., terrorist attacks and police punishments)
- Current sympathy value<sup>48</sup>
- Current competence value

### **2.2.2 Actor Rules**

The actor rules tab of the simulation setup window is displayed in Figure 8 and explained in the following subsections. It enables the user to set each of three sets of model parameters, a set for ordinary citizens, a set for police, and a set for terrorists.

---

<sup>48</sup> This and other terms used in the model are explained in Section 2.2.3.1.



**Figure 8. Actor Rules Setup Window**

### **2.2.2.1 Citizens**

Citizens are the basic actors in the model. The user can specify the number of citizens that are to inhabit the initial landscape. The user can also specify a sympathy reversion factor and a social fluidity factor for these basic actors.

#### **Number of Citizens**

The number of citizens is the initial number of passive citizens (i.e., non-police and non-terrorists) that are to be placed on the landscape. This number, plus the initial number of police and terrorists must not exceed the number of lattice cells available (as determined by the width and height values of the environment).

#### **Sympathy Reversion**

The sympathy reversion factor is the value that each actor's sympathy for the terrorist's cause is decremented at each iteration of a simulation. It is intended to capture the notion



that political beliefs will inevitably weaken over time and tend toward a globally neutral position. The sympathy value is maintained internally as a (real) number in the (closed) interval [0,1].

### **Social Fluidity**

Social fluidity represents the “likelihood of breaking a social relationship.” It is intended to reflect the occasional and inexplicable randomness of broken links in an organization or group defined in terms of links between organizational members.

### **2.2.2.2 Police**

Informally, police are those basic actors of the model who are diametrically opposed to the terrorist cause. Citizens are police actors upon model initialization if their (randomly assigned) sympathy lies within the interval  $[0, <\text{Government Threshold}>^{49}]$ . The user may specify an initial number of police for a simulation. The user may also set a police sympathy effect and a maximum [police] precision.

### **Number of Police**

While it is not necessary to specify an initial number of police actors—citizens may be randomly assigned a sympathy value that makes them, by definition, police, or they may transform into police during a simulation—it is sometimes convenient to have some fixed number of police at the beginning of a run. This parameter can also be used to ensure a certain realistic proportionality of police-to-the-general-population for a simulation.

### **Police Sympathy Effect**


The police sympathy effect reflects how much a citizen’s sympathy for the government decreases after the police punish terrorists. Intuitively, the idea is that all police actions have some, even if a very small, deleterious effect on overall government support by the citizenry.

---

<sup>49</sup> The government threshold value can be set by the user within the advanced options tab of the simulation setup window.

### Maximum Precision

The maximum precision parameter refers to the physical range available for police action. A maximum precision value of two, for instance, indicates that a police actor at the center of the small lattice in the Figure 9 can have an effect only on actors located at one of the shaded squares.



**Figure 9. Police Precision**

### 2.2.2.3 Terrorists

#### Competency

The “competency” of an actor is used to represents the actor’s inherent strength (or power). Actor competency is also represented within the model as a real number in the interval  $[0, 1]$ . The “competency” of a terrorist or counter-terrorist changes as a function of the number of successful terrorist “attacks” or counter-terrorist “actions” in which the actor participates.

#### Type

The “type” of an actor is a function of the actor’s “sympathy” and “competency.”

#### Terrorist Sympathy Effects

The terrorist sympathy effect value is the amount a terrorist’s attack diminishes the sympathy of citizens indirectly affected by the attack.

**Competency Increment**

The competency increment is the increment by which a terrorist's competency is increased by virtue of a successful attack.

**Competency Reversion**

The competency reversion is the competency penalty suffered by terrorists if they do not engage in terrorist activity.

**Maximum Links**

This value is the maximum number of links a terrorist can have with other terrorists in the formation of a terrorist network.

**Minimum Interval**

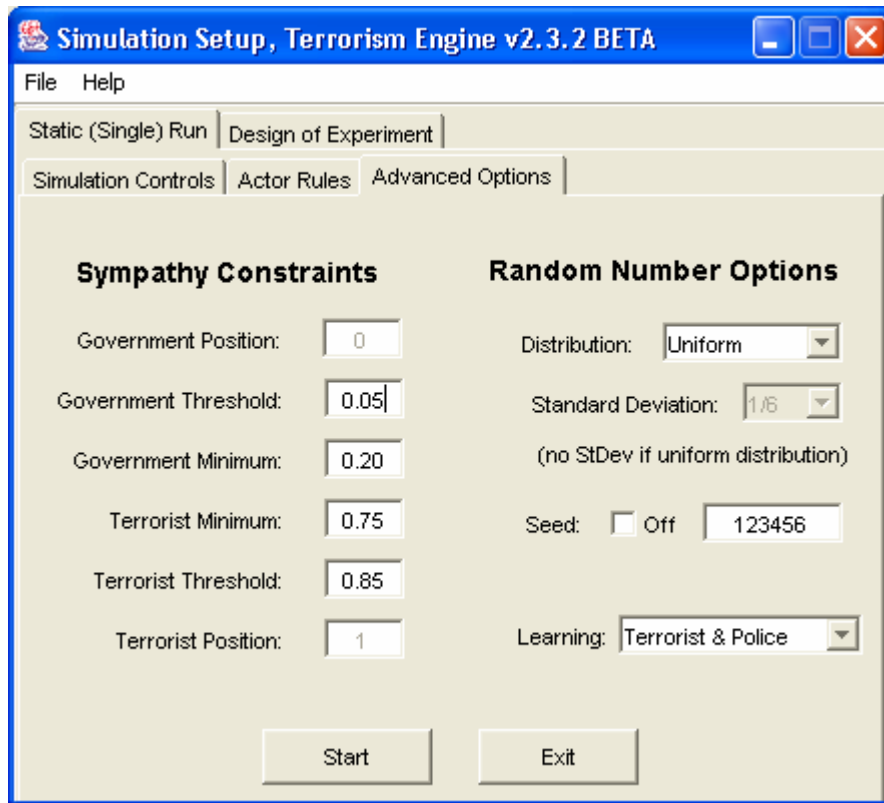
This is the minimum number of iterations that can occur between terrorist attacks.

**Maximum Magnitude**

This is the largest magnitude of a terrorist attack.

**2.2.3 Advanced Options**

There are two sets of advanced options available for the current model. The set of variables that determine an actors fundamental type (passive citizen, police, etc.) is collected under the heading "Sympathy Constraints." A set of parameters then enable the user to control certain randomizing aspects of the model is labeled "Random Number Options." Because agent learning also involves the use of randomizing mechanisms, that control—labeled "learning"—is available within the set of random number options. The advanced options tab of the basic simulation setup window is shown in Figure 10.



**Figure 10. Advanced Options**

### **2.2.3.1 Sympathy Constraints**

Actors have three different attributes: competency, type, and sympathy. The sympathy attribute is manipulated by a set of constraints with respect to both the government (or police) and terrorist positions. This sympathy attribute is intended to represent the actors attraction or aversion with respect to a political position. One could think of this attribute as a value lying on a number line that ranges from  $-1$  to  $+1$ , with  $-1$  representing one extreme position and  $+1$  its diametrical opposing position. A value of  $0$  reflects complete indifference to either position. (For programming convenience, we use the range  $0$  to  $1$ , with  $0.5$  representing the theoretically indifferent mid-point position.) In reality, however, things are never so starkly contrasted, and so we represent actor sympathy “states” (e.g., counter-terrorist, passive citizen, latent terrorist) by partitioning the number line into various intervals. For instance, an actor whose “sympathy” lies in the interval  $[0.25, 0.75]$  might be said (somewhat arbitrarily) to be a passive citizen.

The initial “sympathy” values of the set of actors is either a normal or uniform random distribution. These values then change as an effect of the simple passage of “time” (iterations of a simulation run)—a “tendency” to revert to the theoretical neutral or indifferent position; or in response to the indirect effects resulting from the behaviors of other actors, in particular, terrorist “attacks” or counter-terrorist “actions.” The specifics of this element of the model’s dynamics are fully explained below.

The user has control over these sympathy constraints in the advanced options of the simulation setup window. The specific sympathy constraint options are described below.

#### **Government Position**

The value chosen to denote the government position. Currently it cannot be changed by the user.

#### **Government Threshold**

The government threshold is the point at which citizens become police, if terrorists are sufficiently competent.

#### **Government Minimum**

The government minimum is the point at which police revert to being passive citizens as their sympathy for the police wanes (or their proclivity to the terrorist cause waxes).

#### **Terrorist Minimum**

The terrorist minimum is the point at which terrorists revert to being passive citizens as their sympathy for the terrorist cause wanes (or their proclivity for the government position increases).

#### **Terrorist Threshold**

The terrorist threshold is the point at which citizens become latent terrorists (and susceptible to recruitment by actual terrorists). Latent terrorists become actual terrorist if terrorists are sufficiently competent to recruit them.

### **Terrorist Position**

The value chosen to denote the terrorist position. Currently, it cannot be changed by the user.

### **2.2.3.2 Random Number Options**

Randomization is used extensively throughout the model.

#### **Distribution**

Parameter values are assigned initially in a (pseudo-) random manner using the random number generators available in Sun Microsystems's standard Java<sup>®</sup> system development kit.<sup>50</sup> The user can specify whether the resulting frequency distributions are uniform (i.e., values are distributed approximately evenly over the range of possible values) or normal (i.e., reflecting a more realistic distribution in which the distribution of values are symmetric and have bell-shaped density curves with a single peak). Normal distributions have the nice property that 99.7% of all values lie within three standard deviations of the mean. Different standard deviation values (or different means) yield different normal density curves and hence different normal distributions. Our model currently allows the user to set the standard deviation as either 1/4 (.25) or 1/6 (.16). The standard deviation adjustment is not applicable to uniform distributions. The frequency distribution mode is specifiable within the advanced options tab of the simulation setup window (see Figure 10).

#### **Seed**

For purposes of debugging and more in-depth analysis of simulation results, it is useful to be able to “seed” the random number generator to enable exact repeatability of simulation runs. The advanced options tab of the setup screen allows the user to specify such a seed for the randomization functions. This value is then recorded along with all other initial parameter values of the model in the automated results capture file (see Section 0).

#### **Learning**

---

<sup>50</sup> Specifically J2SE 1.4.2.

The learning pull down menu allows the user to engage the various “learning” algorithms in the model. The user can select either terrorist learning, police learning, terrorist and police learning, or no learning whatsoever (i.e., learning is completely disabled).

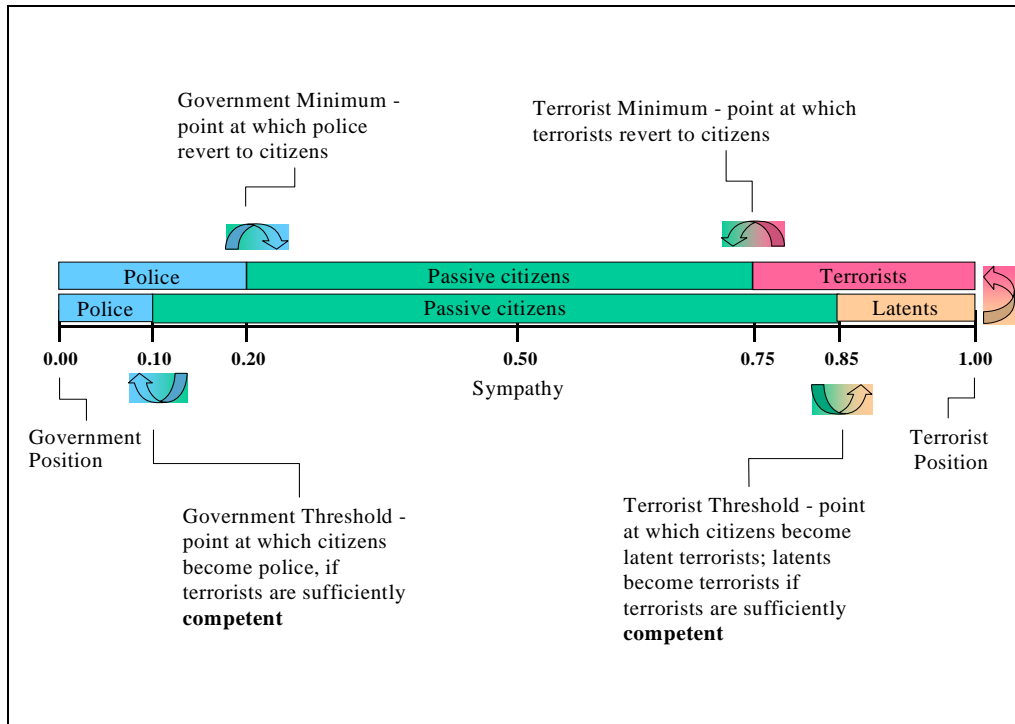
## **2.3 Action Behavior**

There are currently six fundamental behavior rules that govern the behavior of the agents (actors) in the model. All agents are afforded the opportunity to move and transform their sentiment with respect to terrorism or the government position which the terrorists oppose. Terrorists can recruit latent terrorists and form terrorist networks, they can engage in terrorist activities (i.e., attack the authorities or passive citizens), and revise their levels of trust with network members. The police currently can only attack known terrorists.

### **2.3.1 The Sympathy Continuum**

The terrorist behavior model employs one generic kind of actor (or agent). All actors are in one of five dynamic and mutually exclusive states: police (or anti-terrorist), latent anti-terrorist, passive citizen, latent terrorist, and terrorist. Moreover, terrorists, are either known to the counter-terrorists or not. These states are a function of each actor’s “sympathy,” the perceived competence of known terrorists, thresholds, and whether the actor’s “sympathy” is increasing or decreasing. Roughly, a passive citizen becomes a latent terrorist as the actor’s sympathy for the terrorist’s cause exceeds some pre-defined threshold and the known competency level of the terrorist exceeds some pre-defined value. Conversely, a passive citizen becomes a latent counter-terrorist when the actor’s aversion to the terrorist’s cause exceeds some pre-defined threshold. Both terrorists and counter-terrorists can revert to being passive citizens as their “sympathy” for the terrorist or counter-terrorist position changes, coupled with the competency of the terrorist networks. Figure 11 depicts the thresholds at which actors move from one fundamental actor type to another.

Remember that the initial distribution of actor types (counter-terrorist, passive citizen, terrorist) is settable as either normal (Gaussian) or uniform.



**Figure 11. The Sympathy Continuum**

### 2.3.2 Alliances

In addition to the basic behavior rules described above, terrorists and police can form alliances. A terrorist alliance can be thought of as a terrorist network (or organization). An alliance confers additional strength to the individuals that compose the alliances.

### 2.3.3 Adaptation

Both terrorist networks and the authorities adapt to the environment in an attempt to improve their overall performance. Adaptation in the model is implemented by computing two overall utility functions,  $U_p$  and  $U_t$ , for the police and for terrorists, respectively. Both groups then attempt to improve their respective utility value by modifying their rules of behavior. For the terrorists, there are currently three rules they can change: the frequency of attacks, the magnitude of the attack, and the maximum numbers of links that can be maintained in a network. Modification of these key parameters for adaptation purposes is random. This randomness in effect introduces a meta-adaptation principle.

Currently the police can only change the magnitude (precision) of their responses.



### 2.3.4 Innovation

One key design goal of the model was to allow for some form of innovation on the part of the terrorist networks. We felt that innovation was an important aspect of adaptation. Not only would the terrorist networks modify their behavior to improve on the overall outcome of their activities, but they could possibly introduce novel or innovative techniques into their attacks (or, alternatively, devise novel or innovative ways to avoid capture)<sup>51</sup>. (Using fully-fueled, sparsely loaded civilian passenger aircraft as “guided missiles” is one example of a novel and innovative method of attack that was not genuinely anticipated.) The basic idea was to see if an adaptable model of a terrorist organization might not come up with a novel or surprising means of attack in order to improve its overall utility function.

The problem was to determine how such innovation might be introduced into what amounts to a deterministic model. We tried to use the knowledge-base product called Cyc<sup>®</sup> from Cycorp, Inc.,<sup>52</sup> of Austin, Texas, to address this problem.

Cyc<sup>®</sup> represents almost two decades of research and formalization of common sense knowledge and reasoning. The knowledge-base contains over a million rules for reasoning with common sense knowledge, and contains hundreds of thousands of common sense “facts” or assertions about the physical world. The basic means of knowledge representation is a formal language called CycL<sup>®</sup>, an extension of the familiar first-order predicate calculus. Knowledge is organized in a hierarchical network of “microtheories” (or contexts), which comprise a group of related assertions. Cyc<sup>®</sup> contains an inference engine and inference control mechanisms such as forward- and backward-chaining.

---

<sup>51</sup> Many insects are known to have devised remarkable defensive strategies to ensure their survival. Some strains of malarial mosquitoes in Africa “learned” to avoid contact with the DDT sprayed on tent walls by simply flying into the tent, biting the victim, and flying right back out—a “hit-and-run” tactic. If a diamondback moth lands on a tainted leaf, it will fly off but leave its poisoned legs behind—the “leg drop” technique. Some mosquitoes can actually ingest massive doses of certain insecticides, having successfully developed an internal antidote for the poison—the classical “develop-an-immunity” approach. Finally, there’s the “internal-dodge” trick: “The poison has a target somewhere inside the insect’s body. The insect can shrink this target, or move it, or lose it” (Weiner 1994, p. 254f.).

Two versions of Cyc<sup>®</sup> are available. OpenCyc<sup>53</sup> is a free, publicly available, but much smaller version of the complete Cyc<sup>®</sup> knowledge-base called the Integrated Knowledge Base (IKB). In terms of objects (or “things”), the IKB is about one-and-one-half times larger than OpenCyc. There are also approximately two-and-one-half times more micro-theories in the full IKB than there are in OpenCyc. The IKB contains a considerable amount of “knowledge” regarding domains of particular interest to the military, as well as some information on terrorist organizations. Use of the IKB, however, requires a license which would make a widespread dissemination of any model we developed somewhat impractical. Both OpenCyc and the IKB are accessible via a set of Java-based APIs developed by Cycorp. We had access to and employed both the IKB and OpenCyc in our research.

#### **2.3.4.1 Use of the IKB**

The exploratory approach we took initially to using Cycorp’s IKB was to try to take advantage of two inferencing mechanisms afforded by the Java API for CycL<sup>®</sup>: *getGenls* (generalize) and *getSpecs* (specialize). The IKB query *getGenls(Weapon)*, for example, returns a list of all generalizations (e.g., *Individual*, *PhysicalDevice*) of weapon in the IKB knowledge-base. The query *getSpecs(PhysicalDevice)* would then return specific kinds of physical devices (e.g., *RoadVehicle*, *Motorboat*, *HandTool*, *PlumbingFixture*). Beginning with a generic “weapon,” we envisaged a random exploration of IKB concepts related to the concept of a “weapon” by repeatedly invoking the *getGenls* and *getSpecs* API calls.

Our original idea was to use these two inferencing mechanism to essentially “randomly walk” (or explore) the IKB to “equip” terrorists with “novel” kinds of “weapons.” Specifically, *getGenls(Weapon)* would return a list of concepts that were generalizations of “weapon,” for instance {*Individual*, *PhysicalDevice*,...}. This list would then be randomly permuted and the specialization of the first term of the resulting list would be obtained with

---

<sup>52</sup> <http://www.cyc.com/>

<sup>53</sup> <http://www.opencyc.org/>

the *getSpecs* function, for instance, *getSpecs(PhysicalDevice)*. A random selection from the resulting list might suggest the consideration of using a motorboat as a weapon. If the use of these novel (or innovative) “weapons” contributed to greater efficacy of the terrorist networks, we could examine the simulation to see exactly why this novel approach seemed to succeed. Possibly, we could be alerted to the possibility of terrorists using innovative, novel methods and, again, possibly, be able to look for indicators in the real world for the precursors to such behavior. Had the stratagem of using commercial airliners as “guided missiles” under the control of terrorists as pilot occurred, then the idea of looking at flight schools and their students might have prevented the events of September 11, 2001. That, in any event, was the basic idea.

Although we are not ready to give up on this idea completely, we ran into problems that prevented us from fully exploring the feasibility of this approach to genuine innovation in the model. One problem was that the query *getSpecs(getGenls(Weapon))* did not reliably return results (after an indefinite number of iterations). This problem was either due to a software bug in the IKB (or its API) or due to performance problems which we could not resolve.<sup>54</sup>

A more serious problem, however, lay in figuring out just how to incorporate these “novel” techniques (e.g., employing a motorboat as a weapon) within the model itself. The basic question would be how to adjudicate the outcome of an attack of a terrorist group using a motorboat as a weapon against civilians or the police armed with traditional weapons (if armed at all). In terms of the current version of the model, the problem is one of assigning a “capability” value to any actor in possession of a motorboat that would enable adjudication of “attacks” involving a motorboat in a way that was halfway credible. Attributes such as size, weight, inherent dangerousness—as evidenced by the fact that licenses are required for their operation, that there are laws governing their (safe) operation, and that they use flammable and potentially explosive fuels—could certainly figure into such adjudication equations. We did not have the time to explore this avenue to any but this cursory depth.

---

<sup>54</sup> We referred the problem to Cycorp software engineers but have not yet received an adequate response.

#### **2.3.4.2 Use of Open Cycle**

Our use of OpenCyc was basically “proof-of-concept,” intending to demonstrate the feasibility and to assess the utility of using it, initially, merely to enhance the ease with which the model could be modified. There is nothing we did with OpenCyc that could not have been done directly in Java in the model proper. OpenCyc (and especially the IKB) potentially offers access to an enormous amount of information that could not be captured in any practical way within the model itself. We’ve not yet determined how to best make use of this enormous knowledge-base, but proving that it could be easily accessed and used from within the model was worthwhile.

Currently, OpenCyc is used as a knowledge-base to capture salient facts about typical terrorist “weapons,” specifically assassinations, car bombs, and truck bombs. Each of these “weapons” has four (unary) attributes. A competence increment (`competenceIncrement`) is the competence level a terrorist (or terrorist group) must have in order to be able to employ this particular weapon. The values currently in place suggest that the use of a truck bomb requires more competence than that required for a car bomb, which, in turn, requires more wherewithal than an assassination. A sympathy effect (`sympathyEffect`) is the degree to which a passive citizen’s sympathy for the terrorist cause is weakened or lessened when a terrorist uses that particular weapon. A truck bomb has a more deleterious effect than a car bomb which, in turn, has a greater impact on citizen sympathy than does an assassination. The magnitude (`hasMagnitude`) of the weapon captures the scope of the weapon’s effects, again with the truck bomb having greater impact than a car bomb or assassination. Finally, one weapon—the truck bomb—presupposes a certain kind of skill to be used, namely “technical” (`ActorSkill-technical`). These three methods of attack are also considered to be ordered in terms of “technical advance” (`technologyAdvance`) from the lowest level of technology—the assassination—to the highest, the truck bomb.


Actors are randomly assigned certain skills: financial, leadership, and technical. Only actors with “technical” skill can carry out a truck bombing.

## 2.4 Design Experiments

The design of experiments mode is designed to enable the researcher to systematically explore different regions of the parameter space, looking for those parameters or specific parameter values that are in some sense “significant.” The identification of any such significant parameters (or parameter values) may provide insight into the dynamics of the agent behavior exhibited by the model and, by inference, of analogous behavior in the physical world. The basic design of experiments interface window is shown in Figure 12. It provides a pull-down menu that enables the user to select parameter option sets that correspond (roughly<sup>55</sup>) to the three basic sets available from the static (or single run) mode. The idea behind the design of experiment mode is to provide a mechanism by which the user can specify parameters of interest and the range of values over which those values should be methodically invoked during a series of simulation runs. In effect it allows the user to specify, in addition to all the usual initialization parameters, the start-value, increment-value, and end-value (inclusive) for any parameter (or set of parameters) that the user is interested in exploring.

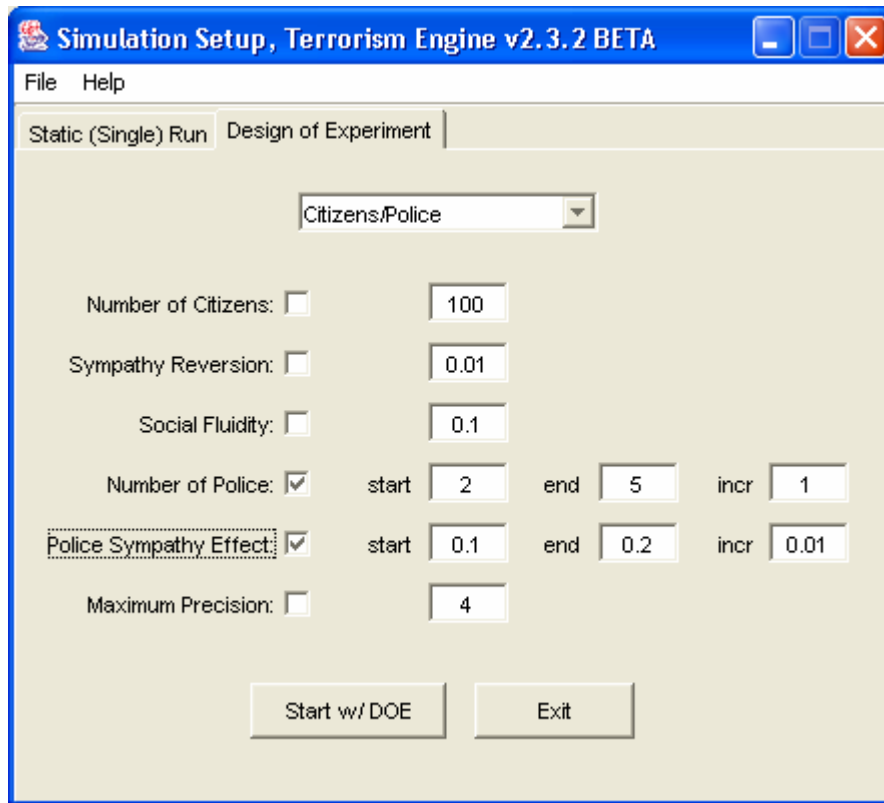
---

<sup>55</sup> There are three primary sets of parameter options available under the Static (Single) Run configuration mode: simulation controls, actor rules, and advanced options. The Design of Experiments mode options map one-to-one to simulation controls and advanced options, but it breaks the “actor rules” tab into two option sets, “citizens/police” and “terrorists.”



**Figure 12. Design of Experiments Interface Window**

In Figure 13, for example, the user has indicated the desire to run the model continuously with different initial values for number of police and police sympathy effect. (Selecting a check-box gives the user additional text-entry boxes in which to specify start, end, and increment (incr) values.) The initial run (consisting of the number of iterations and individual runs at this parameter setting as specified within the environment/output tab) will have two police with an initial police sympathy effect value of 0.1. The second design of experiment run—invoked automatically—will have two police also, but the police sympathy effect value will be at 0.11 (the start value, 0.1, plus the increment, 0.01). Each design of experiment run is an instance of the cross-product of the set of values, {start, start + incr, ..., end} of each parameter for which the user has indicated a design to simulate. In the setup indicated in Figure 13, there will be 44 ( $\{2, 3, 4, 5\} * \{0.1, 0.11, \dots, 0.2\}$ ) separate design of experiment runs.



**Figure 13. Design of Experiments**

The design of experiments mode greatly facilitates both the systematic generation and, when coupled with the various statistical analysis techniques available in Microsoft® Excel, the analysis and graphical presentation of a large number simulations. These results can then be used to guide the design of subsequent simulations in order to eventually converge—it is hoped—on results of genuine value.





### 3. Preliminary Results and Interpretations

---

In this section we report on some of the preliminary results gained from using the model along with some tentative interpretations of these results. Conclusions, both those that might be drawn from these findings and our overall assessment of the value of this approach (i.e., agent-based modeling and simulation) to counter asymmetric threats is provided in Section 0.

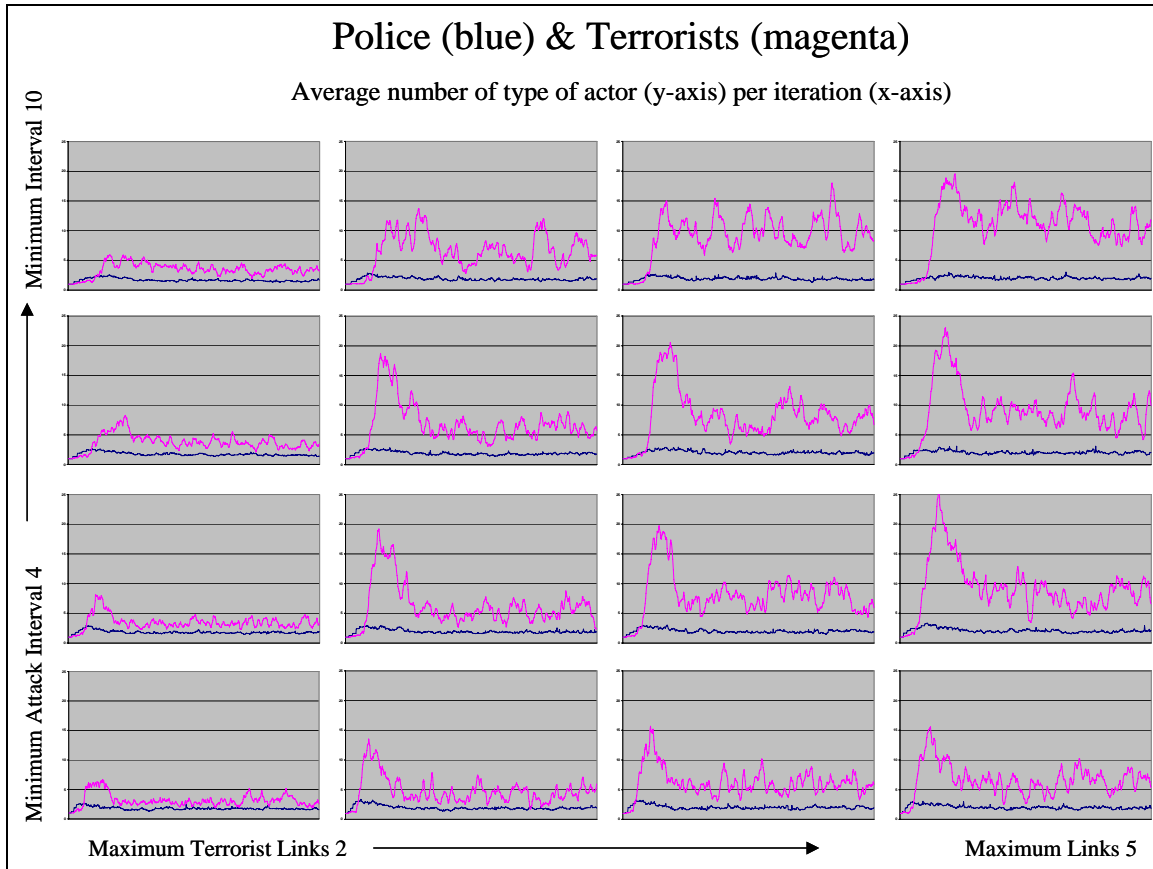
In its current (August 2003) configuration, the model exhibits the effects of three predominant dynamics: ideological fluidity as a function of actor sympathy/aversion and perception of competence toward a terrorist cause, strength as a function of the ability to pool resources by forming networks, and adaptation or the adjustment of methods to maximize an overall utility function. (As noted in Section 2.3) Adaptation (“learning”) is implemented by allowing the variation of three relevant parameters: the scale of a terrorist attack or a police response, the minimum number of members of the organization (police or terrorist) necessary before the organization can mount a response or an attack, and, in the case of terrorist groups, the interval (measured in number of iterations) between attacks.) It seems reasonable, therefore, to examine the behavior of the model with these three viewpoints in mind. In other words, a methodical analysis of the behavior of the model ought to focus (at least initially) on the dynamics of ideological fluidity, the salient factors, if any, of forming alliances or networks, and, finally, the effects of adaptation. Since the latter is, in a sense, orthogonal in nature to the other two perspectives, we should compare the effects of adaptation on the results obtained without adaptation enabled. This methodological approach is then the basis for the organization of this section.

We looked at the effect on the (average) number of police, terrorists, and latent terrorists; the number of events (i.e., terrorist attacks); and the average value of “sympathy” and “competence” as a function of difference parameter settings. These different parameters settings fall into four groups: (1) minimum interval (i.e., number of iterations) between terrorists attacks, (2) magnitude of terrorist attacks and precision of police response, with learning both enabled and disabled, (3) magnitude of terrorist attacks and maximum

number of network links allowed in a terrorist network, again with learning either enabled or disabled, and (4) terrorist sympathy and police sympathy. Table 1 summarizes 17 different configurations for the results that were obtained that are described below.

Input Parameters		Output		
		Average Number of Police and Terrorists	Average Number of Police, Terrorists and Latent Terrorists	Average Number of Terrorist Events, Sympathy (Values), and Competence (Values)
1	Minimum Attack Interval (4→10) vs. Maximum Terrorist Links (2→5)	Figure 14		Figure 15
2	Terrorist Attack Magnitude (2→5) vs. Police Precision (2→5)		Figure 16	Figure 17
	Terrorist Attack Magnitude (3, 5) vs. Police Precision (3), Terrorist Learning Enabled/Disabled		Figure 18	Figure 19
	Terrorist Attack Magnitude (3, 5) vs. Police Precision (5), Terrorist Learning Enabled/Disabled		Figure 20	Figure 21
3	Terrorist Attack Magnitude (3, 4, 5) vs. Maximum Terrorist Links (2, 3), Terrorist Learning Enabled/Disabled		Figure 22	Figure 23
4	Terrorist Sympathy Effect (0.1 → 0.25) vs. Police Sympathy Effect (0.05 → 0.2)	Figure 24		Figure 25

**Table 1. Map of Simulation Result Graphs**




**Figure 14. Number of Police and Terrorists  
as Function of Attack Interval and Terrorist Links**

The series of graphs displayed in Figure 14 illustrates the importance of the number of links allowable relative to the number of intervals (i.e., iterations) that terrorists must wait between attacks.

The minimum interval parameter seems to have little effect vis-à-vis the average number of police and terrorists at each iteration of the simulation, as evidenced by the basic similarity of each row of graphs in the figure.


The factor that seems to have the greatest effect on the outcomes of interest, namely the average number of police and terrorists over the course of the simulation run, is the (maximum) number of links allowed in forming terrorist networks. There is an obvious difference in the overall “shape” of the graphs between those in the first column (where the maximum number of allowable terrorist links is two) and those in columns two, three,

and four. The importance of allowable links relative to attack intervals is due, presumably, to the fact latent terrorists do not become actual terrorists until “recruited” by an actual terrorist. By allowing actual terrorists to recruit (i.e., establish links to) more latent terrorists (who become actual terrorists when a member of a network), the number of actual terrorists can grow rather abruptly.



**Figure 15. Terrorist Events, Sympathy, and Competence as Function of Minimum Attack Interval and Maximum Terrorist Links**

The series of graphs in Figure 15 show dramatically the effects of allowing larger terrorist networks (i.e., a larger number of links per terrorist) in terms of the number of terrorist events. The number of such events jumps significantly when only one additional link is allowed (from the first to the second column). The increase in the number of allowable links seems to have to effect an overall level of terrorist sympathy or competence.




**Figure 16. Number of Police, Terrorists, and Latent Terrorists as Function of Terrorist Attack Magnitude and Police Precision**

Figure 16 shows the varying effects of terrorist attack magnitude (how many actors are adversely affected) and police precision (the extent to which police response is measured and appropriate or indiscriminate and heavy-handed) on the average number of police, terrorists, and latent terrorist. As might be expected, the number of terrorists, reflecting the general population's sympathy (support) for the terrorist's cause, is depressed as the terrorist attack magnitude increases, as shown in the first column of the figure. Moreover, the number of police increases rather dramatically under the same circumstances, as is reasonable.

In the diametrically opposite situation when police precision (analogous to terrorist magnitude) increases (i.e., becomes more indiscriminate), the number of latent terrorist increases dramatically, as indicated in columns two through four. This is consistent with our intuitions that sympathy for the terrorist cause would probably increase as the authorities begin to act more like the terrorists they are trying to defeat.

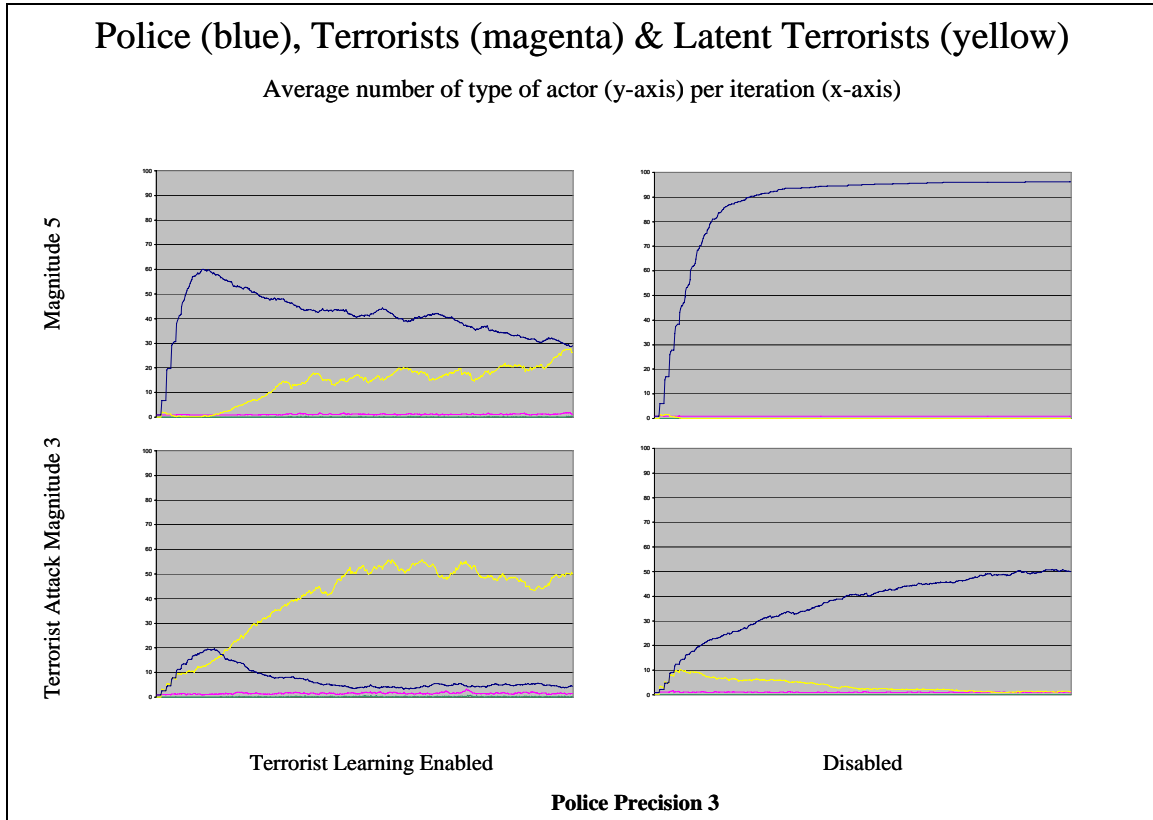
One interesting feature of the figure to note is that in those cases where increases in terrorist attack magnitude and police precision are comparable—the diagonal from lower-left to upper-right—the terrorists seem to be more successful (in terms of numbers of terrorists) than the police. This might be explained by the fact that terrorists learn (adapt) whereas the police do not.



**Figure 17. Terrorist Events, Sympathy, and Competence as Function of Terrorist Attack Magnitude and Police Precision**


Figure 17 depicts the effects on the number of terrorist events, and sympathy and competence levels of the same parameters as Figure 16. Not surprisingly, when terrorist attack magnitude is high and police precision is low (i.e., “good”)—as indicated by the upper-left-most graph—there is very little sympathy for the terrorist cause, a small number of attacks (events) and low competence. The bottom-right-most graph indicates relatively high sympathy for the terrorists and, accordingly, many more attacks.

One should keep in mind, however, that where the number of terrorist attacks appears to be greater in that graphs that appear in the lower rows of the figure, the events graphed along the upper rows affect many more actors or have a greater overall impact.



**Figure 18. Number of Police and Terrorists as Function of Terrorist Attack Magnitude and Police Precision of Three with Learning Enabled and Disabled**

Figure 18 illustrates the effect of learning (adaptation) in the model. With learning disabled, the number of latent terrorists that emerge are well under ten percent of the total population. With learning enabled, however, the percentage doubles when terrorist attack magnitude is five and increases to about 50 percent of the total population at an attack magnitude of three.



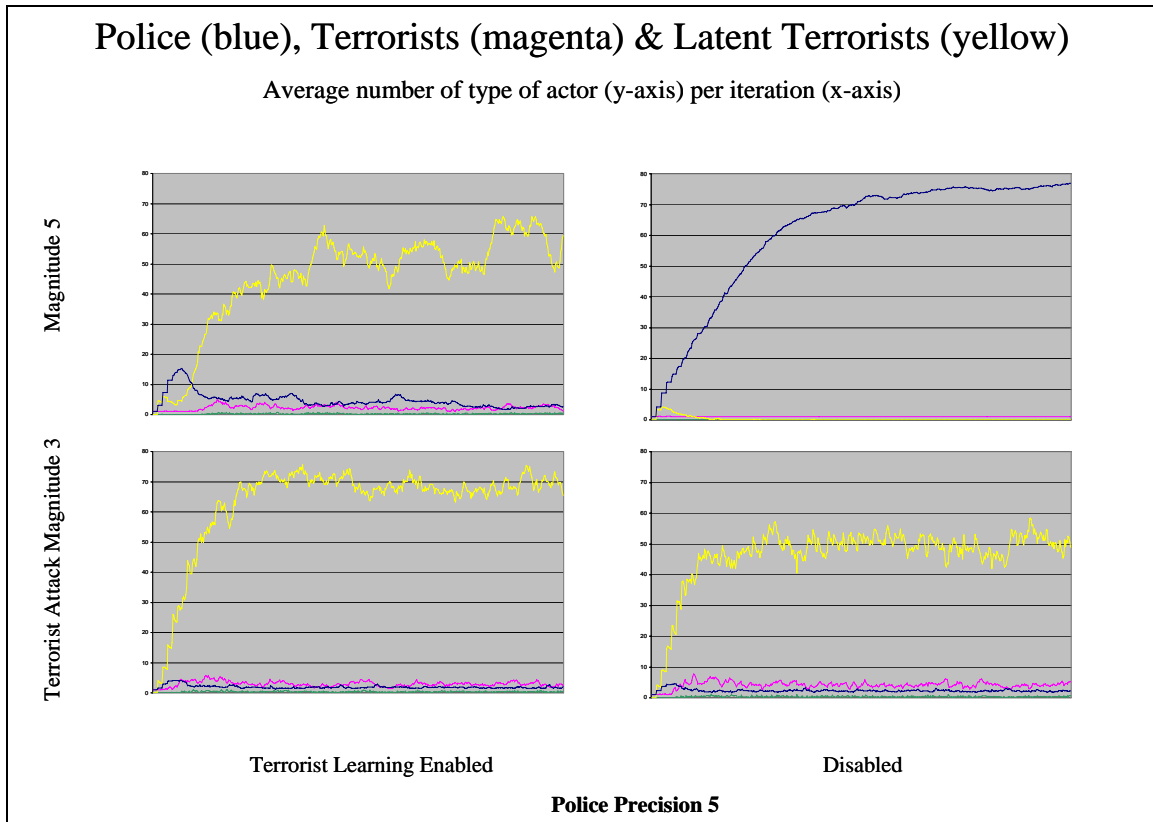
**Figure 19. Terrorist Events, Sympathy, and Competence as Function of Terrorist Attack Magnitude and Police Precision of Three with Learning Enabled and Disabled**

Figure 19 looks at the effect of learning, given a police precision of three, at two levels of terrorist attack magnitude, three and five, of the number of terrorist attacks and global sympathy and competence values.

With learning disabled (the right two graphs), the number of terrorist events trail off very rapidly, especially at an attack magnitude of five. Sympathy also drops below fifty percent and overall terrorist competency falls to about twenty-five percent. The precipitous drop in the number of events is correlated, no doubt, to the equally precipitous drop in terrorist competency.

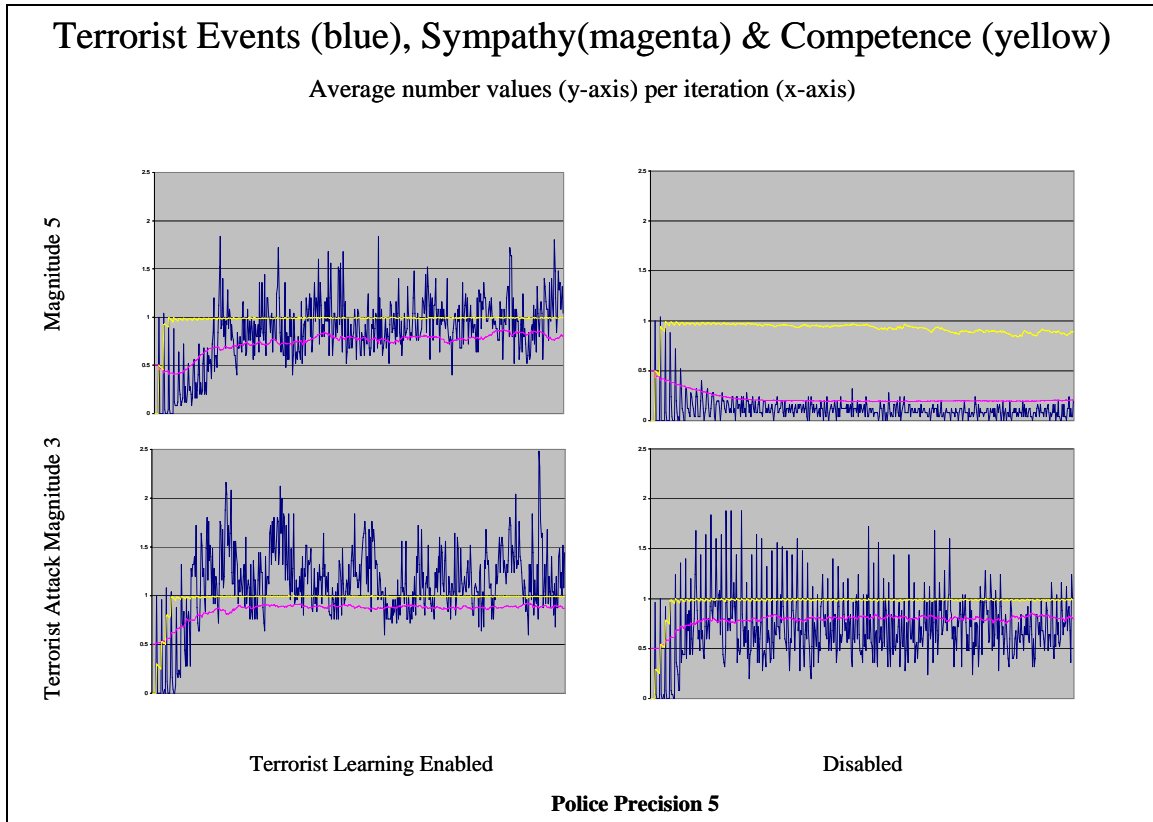
With learning enabled, however, the number of terrorist attacks are more in line with the underlying precipitating conditions. Sympathy hovers around fifty percent at magnitude five and reaches seventy-five percent at a magnitude level of three. Overall terrorist competence remains robust throughout the simulation runs when learning is enabled.





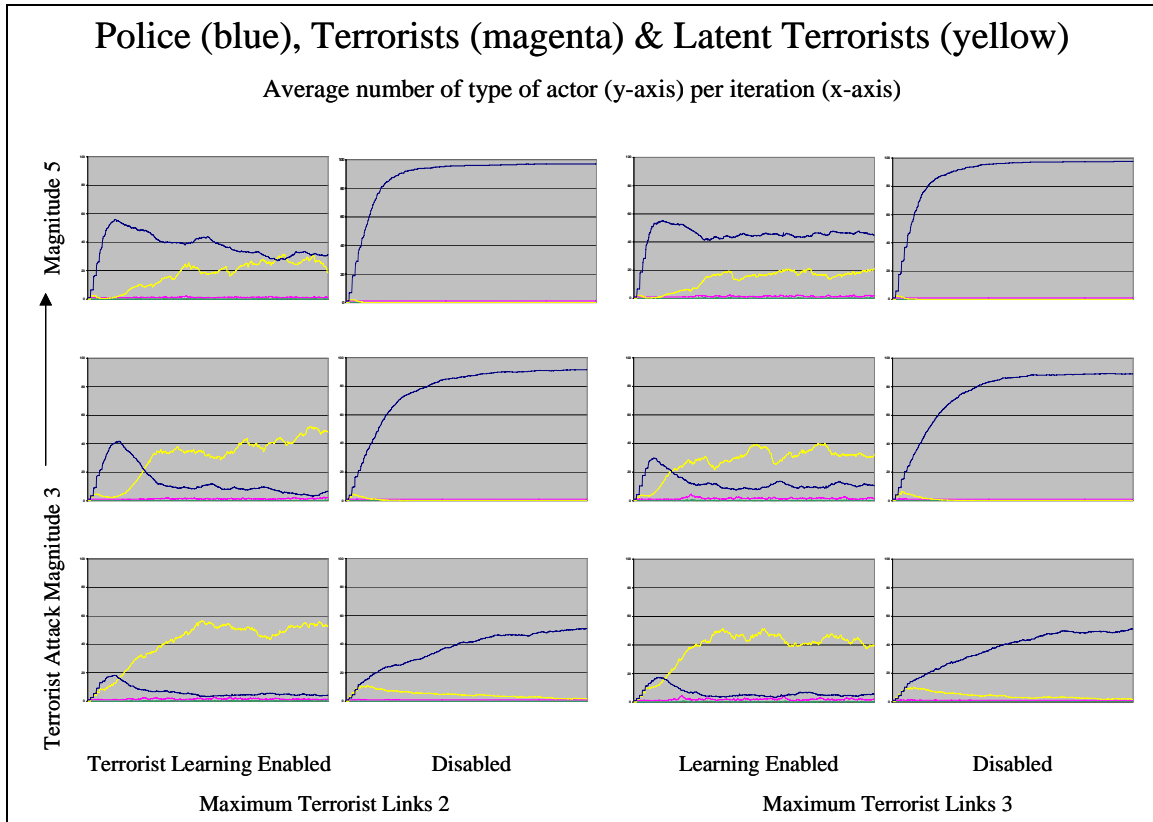
**Figure 20. Number of Police, Terrorists, and Latent Terrorist as Function of Terrorist Attack Magnitude and Police Precision of Five with Learning Enabled and Disabled**

Figure 20 should be looked at along-side Figure 18. Here police precision is set to five whereas in Figure 18 the police precision is three. With learning enabled and a terrorist attack magnitude of five, latent terrorists replace the police in dominating the population. Whereas the average number of police represented over seventy percent of the population when the terrorist attack magnitude was five (see Figure 18), in this case the latent terrorists represent between forty and sixty percent of the population (when terrorist can learn).



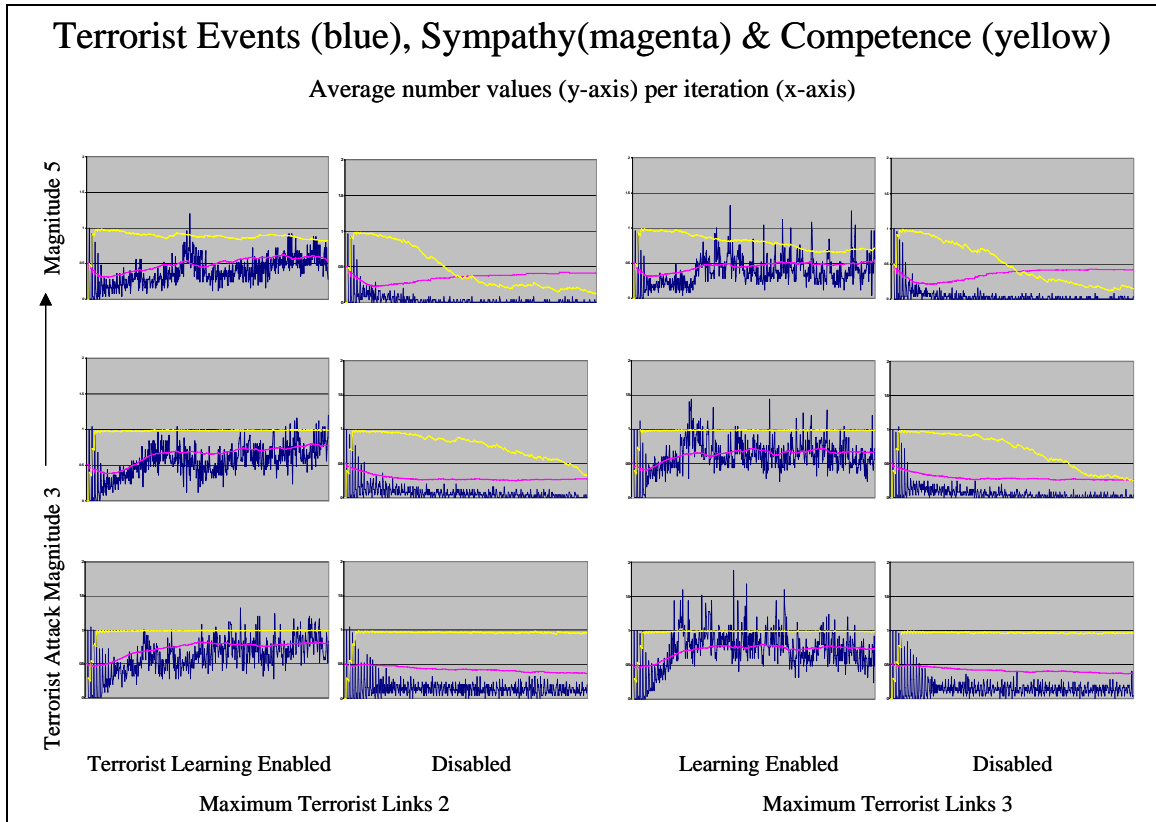
**Figure 21. Events, Sympathy, and Competence as Function of Terrorist Attack Magnitude and Police Precision of Five with Learning Enabled and Disabled**

Figure 21 should be considered along-side Figure 19, particularly with respect to the number of events and the average sympathy values. The upper-left-most graph, especially, with a terrorist attack magnitude of five and with learning enabled, shows a dramatic effect on both the number of events (terrorist attacks) and the average sympathy values.



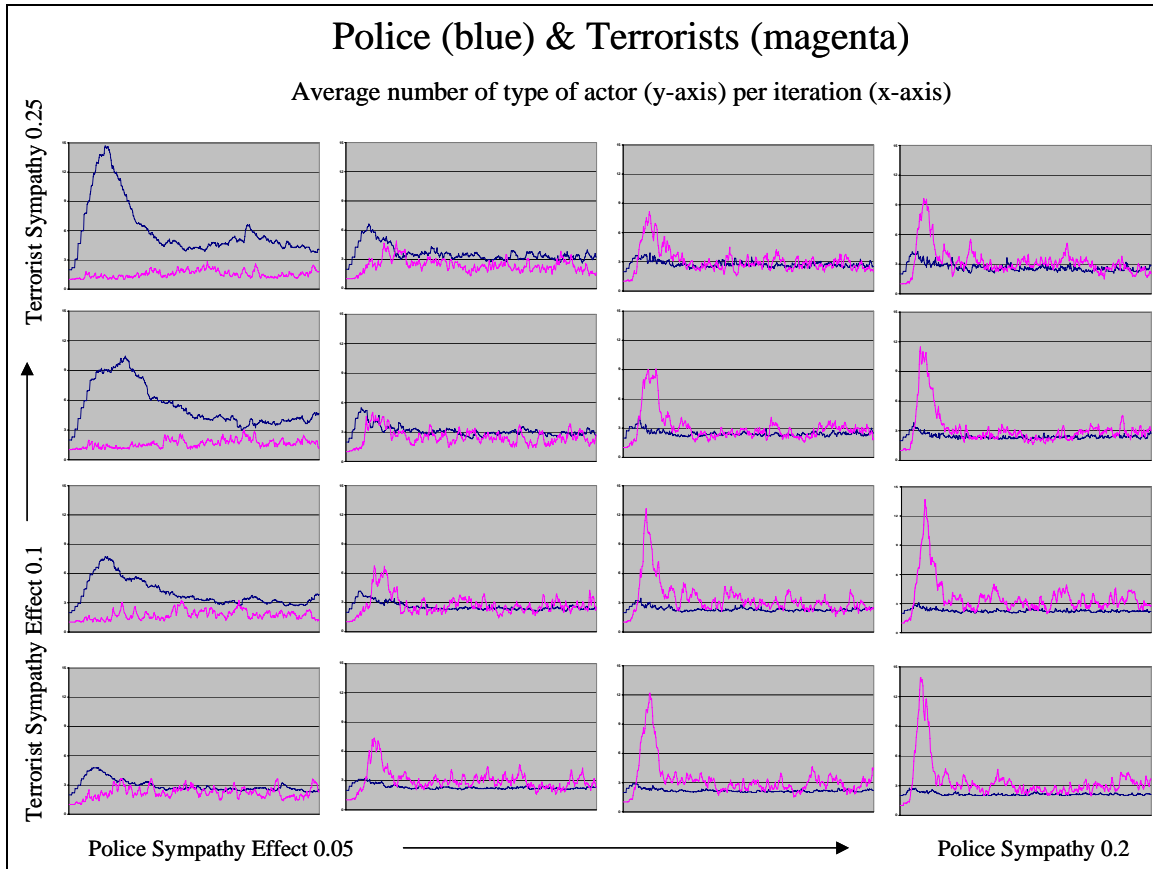
**Figure 22. Number of Police, Terrorists, and Latent Terrorists as Function of Terrorist Magnitude and Maximum Terrorist Links with Learning Enabled and Disabled**

Figure 22 again shows the importance of learning to the terrorist groups. In each set of graphs, the number of police is reduced by about fifty percent when learning is enabled. Moreover, the number of latent terrorists jumps from almost zero to between twenty and fifty percent of the population. The size of the terrorist groups, as reflected in the maximum number of links allowable, seems to have little effect.



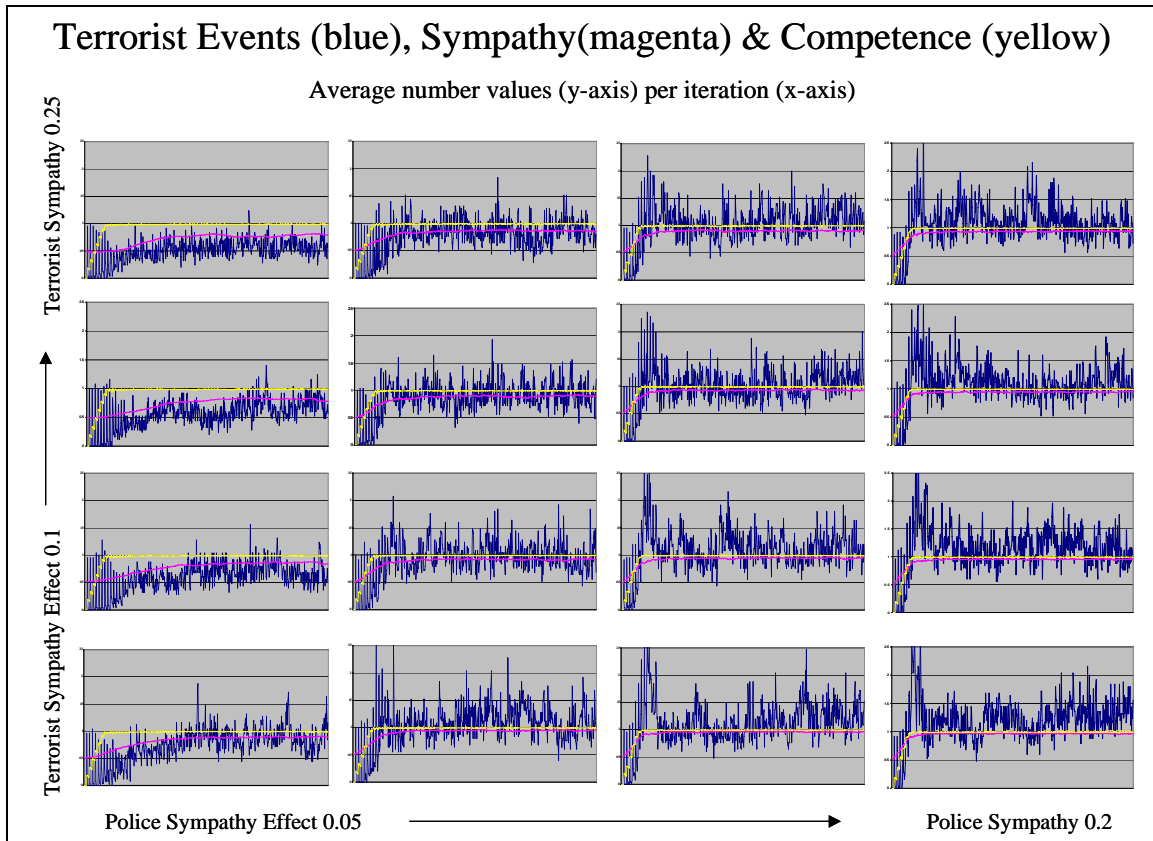
**Figure 23. Number of Events, Sympathy, and Competence as Function of Terrorist Magnitude and Maximum Terrorist Links with Learning Enabled and Disabled**

Figure 23 is another example of the impact of learning on overall model behavior. Terrorist sympathy is less than fifty percent in each of the graphs with learning disabled. It rises to at least fifty percent (and to seventy-five percent) when it is enabled. The number of events (terrorist attacks) jumps as well. Interestingly, the overall perceived level of competence of the terrorist networks is relatively invariant to learning.



**Figure 24. Number of Police and Terrorists as Function of Terrorist Sympathy and Police Sympathy**

Figure 24 shows a series of graphs for varying police sympathy effect and terrorist sympathy effect. (Terrorist sympathy effect is the amount of sympathy for the terrorist cause that a non-political citizen will give up. Police sympathy effect is the amount of sympathy for the police that a non-political citizen will give up.) This particular set of graphs suggests that the effects of terrorist sympathy and police sympathy are more or less equally opposed (or balanced). The set of graphs are practically symmetrical around both diagonals; this is particularly striking when comparing the upper-left-most graph with the lower-right-most graph. This behavior might be consistent with an *ethical relativism* interpretation of real-world behavior: there is nothing inherently different between police action and terrorist attacks; they just reflect the extremes of diametrically opposed behavior.



**Figure 25. Events, Sympathy, and Competence as Function of Terrorist Sympathy and Police Sympathy**

The graphs of Figure 25 show the effect of police sympathy effect and terrorist sympathy effect on the number of events, overall sympathy values, and overall competence perceptions. In contrast to the apparent symmetry of Figure 24, here the police sympathy effect is predominant, especially with respect to the number of events and the sympathy values. It suggests that the authorities have to be more careful than miscreants for fear of alienating a larger percent of the population and encouraging more terrorist activity.

## 4. Conclusions

---

There was one overarching objective to this research—to determine the value of agent-based modeling (and related analytical tools such as social network analysis and systems dynamics) to analyze asymmetric threats (e.g., terrorist groups) as complex adaptive systems. There are two facets to this objective, a tactical facet and a strategic facet. On the tactical side, we wanted to actually apply ABM techniques to a specific problem to see what could be learned about that problem. On the strategic side, we hoped to learn enough about the tools to be able to render a judgment as to their general utility in addressing a certain class of difficult problems. These two facets of the research objective are not independent, but they are not so tightly coupled as to lead us to reject ABM for this kind of problem—a strategic assessment—even if we were to fail in applying the techniques for tactical advantage. It is easy to see how one might fail to derive results that promise immediate practical value yet still reveal how the tools could be applied to significant advantage in the future.

### 4.1 Tactical Value of Agent-Based Modeling

Agent-based modeling is very useful in enabling the “quantification” of causal interactions for which there are no other known and practical techniques.<sup>56</sup> It is not overly difficult to create a model that purports to model complex behavior and attempt to gain insight into physical world complex behavior through synthetic simulation. Simulation behavior, moreover, can be captured and interpreted numerically to render “quantification.” The model described in Section 2 to show how theories of social behavior can be rendered amenable to computational (quantificational) techniques. It does not show us that the results obtained—the answers derived—are of real value (i.e., that the “answers” are correct). The results obtained, to be of genuine value, need to have been derived from a

---

<sup>56</sup> The differential equation approach is theoretically sound but not practical. Simple equations may be solvable but are not realistic; on the other hand, the sheer number and complexity of the equation set presumed necessary to model physical world social behavior are intractable.

“validated” model, and we’ve made no attempt to validate this model of terrorist behavior.

#### **4.1.1 The Need to Validate Agent-Based Models**

There are two steps to model validation (corresponding roughly to the conventional verification and validation (“V&V”) process used in computer software evaluation): (1) Does the model perform internally as it was designed to perform? (2) Does the internal performance correctly mirror the physical world processes the model is intended to model? The first question addresses the correctness of the algorithms that control the model’s behavior. The second question addresses the soundness of the algorithms—do they reflect the way the physical world behaves?—and the reasonableness of the initial conditions of any given simulation—do they represent a plausible abstraction from a real (or possible) physical world situation? This second step of the validation process is a form of calibration. The actual behavior of the model is calibrated—if possible—to accord with behavior in the physical world.

Validation of the correct functioning of the model itself is relatively unproblematic, although it is often difficult in practice: due to “combinatorial explosion,” you simply can’t test everything. Validation against the physical world—calibration—is more difficult, especially if the behavior against which the model is to be compared is not readily available. It’s one thing to calibrate models of vehicular traffic patterns against known traffic patterns as determined from empirical studies (in the form of manual or automated traffic surveys). It’s another to attempt to calibrate a model of behavior that is illegal against the corresponding physical world activities that are largely covert. There is simply too little known about actual activities to which the model’s behavior can be compared in order to validate it. Terrorist behavior, in particular, is notoriously difficult to fathom. The number of terrorist incidents are too few and too heterogeneous to serve as a basis from which to infer generalized behavior patterns. Interrogation of detained terrorists themselves as to



their motives and *modus operandi* is generally regarded as of questionable validity into their true motives and strategies.<sup>57</sup>

Although this suggests that ABM as applied to certain domains (those characterized by essentially covert activities not directly observable) perhaps should not be held to the same standards of evaluation as when the technique is used in more transparent domains, it does not tell us if there are not more applicable standards. Nor does it tell us what those more applicable standards might be, if there are any.

#### **4.2 Strategic Value of Agent-Based Modeling/Simulations Research**

As hinted at previously, we are hesitant to unequivocally embrace the information technology that is the subject of this CRP for near-term, practical value to DoD vis-à-vis certain asymmetric threats (such as transnational terrorism). This is not to say, of course, that neither the research nor the technology is of no or little value. We hope that the research—if sound—sketches the current limits of the value that this particular technology affords to the United States defense establishment. That, in itself, is of considerable value (although tempered, of course, by that caveat “if sound”). Senior DoD decision makers need the conceptual understanding to validly critique the claims made for ABM technology. This is one of the goals of this paper.

The behavior of the simulations these agent-based models inform are given “interpretations” vis-à-vis possible or probably behavior in the physical world that are too often tenuous at best and misleading and dangerous at worst. This cannot be proved. All we can hope to do is look carefully at what’s going on and point out the (often) tenuous interpretations of computer simulations with the physical world phenomena these simulations purport to represent.

The best word of advice we can offer the reader who examines the results of various ABM experiments is to note the abundant use (at least initially) of scare quotes (“”) signaling honestly to the reader that “X” doesn’t really mean X, but could, if one were chari-

---

<sup>57</sup> See fn. 7.

table enough, think of “X” as X. The problem, of course, is that these scare quotes begin to clutter-up the narrative and are consequently dropped with the reader perhaps beginning to believe that, indeed, the behavior of agents *in silico* bear a passing or even a convincing resemblance to agents in the physical world. All of the assumptions—the model’s global parameters—are conveniently forgotten and the trusting decision maker begins to think that here is, indeed, a possible way to foresee the future.

Thus our conclusions are equivocal and non-decisive. While ABM technology may offer significant value in many fields, it is not clear that the technology offers *tactical* value in countering immediate asymmetric threats such as terrorism. The report’s immediate value is in delineating the important questions that should be raised regarding efforts to use ABM technology for immediate operational gain.

We have developed an ABM simulation of terrorist network behavior for the purpose of assessing the usefulness of ABM technology in countering the asymmetric threats imposed by transnational terrorist organizations. The underlying assumption guiding this research is that terrorist organizations (“terrorist networks” is the *nom de jour*) are essentially adaptive complex systems: they exhibit complex behavior in the sense sketched above and they adapt to changing circumstances by modifying their rules of behavior. Accordingly, since ABM techniques were developed in part to analyze complex and arguably adaptive behavior, it seems appropriate to use ABM technology to analyze terrorist networks. Note that to define a terrorist network as a complex adaptive system is not necessarily to distinguish it from any other non-trivial organization. Successful legitimate organizations and institutions of all types can also be considered complex adaptive systems and, hence, amenable to analysis using ABM techniques. Terrorist organizations differ from these legitimate organizations in at least two important respects: their means for accomplishing their goals are different—they use terrorism to further their interests—and, as a consequence, are illicit and subject to law enforcement efforts in ways that legitimate organizations are not. Terrorist organizations can be distinguished from other illicit organizations such as drug cartels by the nature of their fundamental goals. A drug cartel is in business to make money, doing so by illegal means. A terrorist organization uses terrorism to promote a political agenda. Terrorism itself is generally defined as the commis-

sion or threat of acts, usually against innocents or non-combatants, designed to create extreme fear in those who directly or indirectly bring about the changes the terrorist desires. In short, terrorism is the conduct (or threat) of terrorist acts for political purposes. Any ABM simulation that purports to mirror physical world terrorist behavior has to exhibit this defining characteristic in some form or other.

We also insist that the ABM simulation represent terrorist organizations (or networks) as organizations (or networks), that is, as comprised of a group of individuals related to each other (and possibly to others not in the group) in some structured way. Although there are certainly individual terrorists, individuals acting essentially alone are not the focus of our model. (It should not be forgotten, however, that an individual terrorist—one not affiliated with a terrorist network—will probably still need the resources needed by larger, more organized groups. It may be fruitful to think of the individual terrorist, acting alone, as a microcosm of the terrorist group.) At the point of departure, we assume that terrorist networks require resources and have a certain minimal organizational structure. Functionally, a terrorist organization must manage, recruit, train, finance its operations, communicate internally and externally, and operate (i.e., attempt acts of terrorism). They require leadership (a strategist or planner), recruiters, trainers, financial resources, communications facilities, operatives, and so forth. Exactly how terrorist networks constitute themselves in terms of these functions is generally not known, and even when known, such knowledge is of no certain value. We assume that the network, being adaptive, may re-organize repeatedly to better achieve its objectives or to avoid capture or disruption by the authorities.

### **4.3 The Hidden Costs of Agent-Based Modeling**

It would not be unreasonable to say that at least 50 percent (and possibly three-quarters) of the funding expended on this CRP effort was consumed in dealing with mechanics of computer programming and related issues (downloading and installing software, configuring development environments, coding, debugging, testing, etc.). We simply could not find a robust, general purpose, easy to learn and use development environment that would enable “subject matter experts” to specific model behavior in a simple, practical, and effi-

cient way. We came to rely, ultimately, on a development environment that was basically a general purpose Java<sup>58</sup> application programming environment.<sup>59</sup> Building a nice user interface onto the underlying model enabled the research team to easily set the parameters that controlled the behavior of the model. But even then there is, we feel, a considerable “impedance mismatch” between the complexity of the phenomena to be modeled and the technical limits of the tools being used to attempt that modeling.

#### **4.4 The Ultimate Value of Agent-Based Modeling**

The value of agent-based modeling (and simulation) is ultimately, we submit, the value of computer simulation in general. Computer simulations

add...discipline: a way of discovering hidden assumptions of one's models, and a way of exploring the dynamic effects, by “turning the knobs” to see the effect of different settings of the variables. It is important to recognize that these computer simulations are actually philosophical thought experiments, intuition pumps, not empirical experiments. They systematically explore the implications of sets of assumptions. Philosophers used to have to conduct their thought experiments by hand, one at a time. How they can conduct thousands of variations in an hour, a good way of checking to make sure that the intuitions they pump are not artifacts of some arbitrary feature of the scenario. (Dennett 2003, p. 218)

There's immense value in this assumption checking and knob tweaking, but extreme care must be exercised when transferring the insights gained to the “empirical” world. Ultimately a computer-based simulation is the exercise of a computer program, elements of which may or may not be mappable to the physical world. The dots running around on the screen (or lattice cells changing color) are ultimately just that: dots running around on the screen. It is the structure that determines that “running around” behavior that may be the same structure that determines some real world behavior in which we're vitally interested. And in that fact lies the real value and importance of modeling building and simulation. The computer gives us a way to bring a rigorous discipline to our model building and simulation, just as it gives us better tools for exploring the assumptions upon which

---

<sup>58</sup> Specifically, Java 2 Standard Edition (J2SE) 1.4.2.

<sup>59</sup> JBuilder SE (7.0.155.0) from Borland®.

these models are based. It allow us to explore the large possibility spaces<sup>60</sup> engendered by many different possible agent actions, governed by different agent beliefs, attitudes, goals, and dispositions. It cannot, however, relieve us of the obligation to argue for the strength of our assumptions independently of our computer-based modeling. That's something we must do for ourselves.

---

<sup>60</sup> When talking about “possibility spaces,” it’s helpful to keep in mind that there are at least four different kinds of possibilities: logical, physical, biological, and historical (Dennett 1995, p. 104). Logical possibility includes what is physically possible, which includes what is biologically possible, which includes, in turn, what is historically possible. The historical possible includes what is actual. The behavior we are attempting to model with agent-based systems fall mainly in the realm of “historical” possibility, viewed retrospectively from some point in the (probably) near future.



## Appendix A. Terrorism CRPMt

---

Microtheory : TerrorismCRPMt

Bookkeeping Assertions :

(myCreator TerrorismCRPMt CycAdministrator) in BookkeepingMt  
(myCreationTime TerrorismCRPMt 20030822) in BookkeepingMt  
(myCreationSecond TerrorismCRPMt 143923) in BookkeepingMt

GAF Arg : 1

Mt : UniversalVocabularyMt

isa : Microtheory

genlMt : BaseKB

Microtheory Contents :

(competenceIncrement TerroristAssassination 0.1)  
(competenceIncrement CarBomb 0.15)  
(competenceIncrement TruckBomb 0.2)  
(sympathyEffect TruckBomb 0.15)  
(sympathyEffect CarBomb 0.1)  
(sympathyEffect TerroristAssassination 0.05)  
(hasMagnitude TerroristAssassination 2)  
(hasMagnitude CarBomb 3)  
(hasMagnitude TruckBomb 5)  
(requiresSkillToUse TruckBomb ActorSkill-technical)  
(technologyAdvance CarBomb TruckBomb)  
(technologyAdvance TerroristAssassination CarBomb)  
(comment technologyAdvance "The first term is the less advanced technology, the second term is the more advanced")  
(argIsa hasMagnitude 2 RealNumber)  
(arg2Isa hasMagnitude RealNumber)  
(argIsa competenceIncrement 2 Number-General)  
(arg2Isa competenceIncrement Number-General)  
(argIsa sympathyEffect 2 Number-General)  
(arg2Isa sympathyEffect Number-General)  
(argIsa sympathyEffect 1 TerroristWeapon)  
(arg1Isa sympathyEffect TerroristWeapon)  
(argIsa competenceIncrement 1 TerroristWeapon)  
(arg1Isa competenceIncrement TerroristWeapon)  
(argIsa hasMagnitude 1 TerroristWeapon)  
(arg1Isa hasMagnitude TerroristWeapon)  
(argIsa technologyAdvance 2 TerroristWeapon)  
(arg2Isa technologyAdvance TerroristWeapon)

(argIsa technologyAdvance 1 TerroristWeapon)  
(arg1Isa technologyAdvance TerroristWeapon)  
(isa TerroristAssassination TerroristWeapon)  
(isa TruckBomb TerroristWeapon)  
(isa CarBomb TerroristWeapon)  
(isa ActorSkill-financial ActorSkill)  
(isa ActorSkill-leadership ActorSkill)  
(isa ActorSkill-technical ActorSkill)  
(argIsa requiresSkillToUse 1 TerroristWeapon)  
(arg1Isa requiresSkillToUse TerroristWeapon)  
(argIsa requiresSkillToUse 2 ActorSkill)  
(arg2Isa requiresSkillToUse ActorSkill)  
(argIsa requiresSkillToResearch 2 ActorSkill)  
(arg2Isa requiresSkillToResearch ActorSkill)  
(argIsa requiresSkillToResearch 1 TerroristWeapon)  
(arg1Isa requiresSkillToResearch TerroristWeapon)



## **Appendix B. An Analysis of Some Key Assumptions of the Terrorism Model By Matthew C. MacArthur**

---

This analysis is intended to do three things: to render some key modeling assumptions explicit; to evaluate whether or not these choices were wise, logical, or theoretically defensible; and – where necessary – to suggest alternatives. By doing these things, I hope to help you clarify whether this specific Java formalization of terrorism matches your “mental model” of terrorism, or whether it even meets your standards of “sensibility”. I also hope to help you anticipate critics’ challenges to the choices we made in the formalization of this model.

Clearly, I do not discuss all of the assumptions that have gone into the model. I focus specifically on those likely to raise the eyebrows of people who read about or use the model, and those that – if modified or abandoned – could make the model even more interesting.

Originally, I planned to step through the model one class at a time, but it soon became clear that this would produce a rather disjointed and unhelpful document. I now think a clearer approach is to discuss one modeling assumption at a time, and to note the relevant classes in each instance (so you know where modifications would have to be made if you decided to modify the assumption).

### **Sequence of Play**

#### **Relevant Classes: Engine**

Engine controls the flow of the simulation. The sequence of simulation stages is Update Confidants, Transform, Network, Terrorize (only terrorists), Police (only police), and Move. In each stage, all actors who can “legally” participate do so before the simulation

proceeds to the next stage.<sup>61</sup> That is, there is an outer stage loop in which there is a nested inner actor loop (stage-actor).

An alternative would be to nest the stage loop within the actor loop (actor-stage). Indeed, early versions of this model reflected this alternative flow.

I'm mostly agnostic about this difference. In models with few stages, I doubt stage-actor and actor-stage implementations of the same model would produce different results.

However, as the number of stages increases in an actor-stage model, the more activities each actor can accomplish before another actor gets its turn. This can confer very high first-mover advantages to the first few actors lucky enough to stand in the front of the queue. Actions of these actors may turn out to be excessively decisive with respect to the final results of simulation runs. Therefore, I suggest retaining the stage-actor flow structure.

Likewise, I think the flow of stages should be kept more or less as-is. The Move stage could be inserted almost anywhere in the sequence of stages without producing significantly different results. The order of the other stages just seems to be “sensible”.

My biggest concern is that Terrorize and Police are treated as separate stages, and as a consequence all the terrorist actors have opportunities to act before any police actors have a chance to subdue them. In the simulation, police can identify known terrorists by their actions before these terrorists see them, but the police still cannot reach them to arrest them (because their radii of awareness is 3 and their radii of action is only 2). Even known terrorists will *always* have a chance to act just before they are captured – police can never prevent them.

This arrangement seems somewhat unreasonable, so I suggest merging these two stages into a single generic “Action” stage in which terrorists terrorize and police. I would also

---

<sup>61</sup> In some stages, not all actors can participate. For instance, police actors never participate in the Terrorize stage.

consider boosting the police radii of action to equal their radii of awareness (that is, make both 3). I doubt these change will substantially alter the simulation results (aside from slightly retarding the absolute level of terrorist activity), but from a modeling standpoint it seems more “aesthetically correct”, and it may preempt criticism of the model’s design.

### **All Police Are Overt**

#### **Relevant Classes: Terrorism Behavior Set and Police Organization**

An assumption that might be fruitfully abandoned is that all police are overt. You should consider introducing undercover police or informers who (unlike overt police in the model now) can enter into confidant relationships with other actors. These actors can then “report” actors who attempt to recruit them into terrorist organizations. A more sophisticated approach would be to allow police organizations to have different covert strategies – they may act as soon as they discover the identity of one terrorist, or they may lie low until they have uncovered some specified number of terrorists.<sup>62</sup>

This would obviously represent a major change to the model, but it would open up some new avenues of inquiry. For example, it might be possible to see what the optimal mix of overt and covert police elements is versus a given terrorist organization. An analyst may also be able to trace the effects of covert policing on terrorist organization recruiting activities.

### **Confidants**

#### **Relevant Classes: Terrorist Brain and Terrorism Behavior Set**

The mechanisms governing social association among the non-police actors in the simulation seem adequate for present purposes. After all, this is a simulation of terrorism, not of all social life, so relying on a very simplistic depiction of the process of acquiring and discarding friends (or confidants, in the model) seems forgivable.

To increase the realism of the acquisition of confidants, the model could be changed such that the first confidant of each actor is chosen randomly, but then actors become confidants of their confidants' confidants (a process that seems more realistic). If someone wants a new friend, they look to friends of their friends first. If they have no friends to begin with, they choose randomly. Whether this change would affect the model results substantially, I do not know. It may impact terrorist organization recruiting, because it's possible this change would result in semi-discrete clusters of confidants (cliques). Once a terrorist organization has recruited as many willing members as it can from a clique, it may have trouble penetrating other cliques.

However, the biggest potential problem with the confidant acquisition mechanism is that the assumptions of a homogenous fixed maximum number of confidants and probability of confidant turnover (social fluidity) seems unrealistic. Dispensing with it may prove fruitful. Right now the number of confidants any actor can have at one time is arbitrarily set to a limit of 8. It seems more realistic that actors vary in their "sociability". Reclusive actors may only have 1 confidant (or even none). Open or gregarious actors might be able to maintain dozens of confidants. Sociability is already a Brain-level attribute, so the only necessary code change would be to randomize the assignment of this upper limit on confidants. Likewise, social fluidity could be made to vary across actors by making this a Brain-level attribute, too.

If one also introduced recruitment strategies to terrorist organizations, it would be possible to see what kinds of people terrorist organizations seek out.<sup>63</sup> Is it better to recruit outcasts or highly-networked people? Furthermore, the terrorist organization maybe granted the ability to dictate the sociability of their recruits (either by increasing or decreasing their upper limit on confidants). One could also examine the interplay between police strategies (particularly with regard to covert operations) and terrorist recruitment strategies.

---

<sup>62</sup> See also the discussion "Known Terrorists" below, as it has obvious implications for the changes suggested here.

## **Uniform Assessment of Competence**

### **Relevant Classes: TerroristBrain, Engine, TerroristBehaviorSet**

One of the reasons I started modeling internal conflict (like rebellion and terrorism) was that I wanted to improve on the work done by Epstein and others in the area. An assumption in this work that I objected to is that all actors regard the legitimacy of their government the same way. More interesting, it seems to me, are situations in which different people have different views of their government.

Ironically, the terrorism model still retains an assumption very similar to the one I found problematic in this earlier work. In the terrorism model, all actors evaluate the competence of the terrorist organization the same way.<sup>64</sup> Given a terrorist action, all the actors adjust their assessments of the terrorist organization's competence the same way.

The clear fix to this problem would be to make the competence increment and decrement variables Brain-level attributes, such that upward and downward assessments vary across actors (though the magnitude of the action should still be one of the determinants of the size of the increments). For some actors, a small action might prove sufficient to convince her of the terrorist organization's competence (high increment). In contrast, other actors might be tougher judges of competence (low increment). Actors may also vary in their interpretation of long periods of terrorist inaction – are they just laying low but retaining competence (low decrement) or are failing to act out of weakness (high decrement)?

### **Known Terrorists**

### **Relevant Classes: Terrorist Brain, Terrorist Behavior Set**

---

<sup>63</sup> Obviously this would require additional changes to the TerroristOrganization class.

<sup>64</sup> I see that competence effects are now connected to the terrorists' chosen weapon. This is an interesting choice, but it still suffers from the same problem I'm describing here.

The problem here is that when a terrorist is known, she *knows* she's known. Many informational issues are assumed away as a result. This assumption also prevents analysis of more sophisticated police strategies by which they gather some level of information on their targets before acting (and “compromising” their sources).

Enriching the model such that terrorists may or may not know they have been “made” would require substantial code changes. First, a new actor state may need to be introduced – something like “known-to-be-known”.

Second, the various behaviors associated with being “known” would have to be re-examined. For example, right now known terrorists do not recruit. This seems like a reasonable assumption – better not to undermine new recruits by linking them to actors already targeted by the police.<sup>65</sup> But if we assumed known terrorists do not know they are known, we should allow them to recruit and should instead restrict recruitment by known-to-be-known terrorists.

Third, changes on the police side of the ledger may also be necessary. How do terrorists ever know they are known? It might be wise to introduce a police organization strategy parameter governing whether or not to put a terrorist on a public “most wanted” list. The police in some cases may decide it is better to keep their information private, in other cases the opposite may prove true. Terrorists would learn they are known if they were named publicly this way. Terrorists would also learn they are known if their recruiters are captured.<sup>66</sup>

## **Sympathy Dynamics**

### **Relevant Classes: Terrorist Brain, Terrorist Behavior Set**

---

<sup>65</sup> Better still would be to make this a parameter of the terrorist organization's strategy. Let the terrorists themselves decide whether restricting recruitment in this way is in their interest or not.

<sup>66</sup> One could further complicate matters by allowing actors to make judgments about whether their comrades will “rat them out or not”. They may not automatically assume that their comrades will reveal their identities, as they do in the model now.

The assumptions regarding increases and decreasing in sympathy toward the terrorists are generally sound, but they do suffer some of the same problems that plague estimation of terrorists' competence. That is, all actors respond basically the same way to various events that affect sympathy. Again, the remedy is a familiar one: make sympathy increments and decrements Brain-level attributes. This would allow for a richer array of actor types: pro-government types (low increment, high decrement), pro-terrorist types (high increment, low decrement), and ambivalent types (roughly equal increment and decrement sizes).

Such a change may yield interesting simulation results. I would speculate that this modification would create a situation of diminishing marginal sympathy benefits resulting from indiscriminate police crackdowns. There would be hard-cases whose sympathies the terrorists could never sway in their favor (the pro-government types described above). Therefore, the terrorist organization will reach a limit where further recruiting is difficult or impossible, and in this circumstance police can operate with a freer hand, unencumbered by fears of driving the population into the arms of the terrorists. Again, this is purely speculation, but clearly this simple change could produce unexpected and possibly fascinating results.

It is also worth noting that competition for sympathy between the police (government) and the terrorists is a zero-sum game. This seems like the correct way to go – especially when one considers how odd it would be to imagine actors somehow sympathetic to both sides – but this assumption should be explicitly noted in any documentation describing the operation of the terrorist model.

### **Only One Terrorist Organization**

#### **Relevant Classes: Probably All Classes**

Critics might suggest that we have neglected the possibility of multiple terrorist organizations. This is true, but it was done purposefully. Introducing multiple terrorist organizations may allow for some very interesting analyses, but it vastly complicates not only the code but also the interpretation of the model's results.

The primary obstacle is that the one-government-versus-one-terrorist-organization model is built on an implicit mono-dimensional policy (or sympathy) space. As discussed above, competition between the police and the terrorists is a zero-sum game. The situation becomes non-zero-sum in nature once one or more additional terrorist organizations are added. To allow for this more complicated competition structure, the code would have to be substantially overhauled.

The second question is how do new organizations form? One way would be to start with (or allow for the emergence of) multiple instigators who build different organizations as they attract more recruits. Another would be to allow for organizational division such that one terrorist organization could spawn “spin-offs” of disaffected members or terrorist “entrepreneurs”.

This change would obviously require a great deal of work, but the benefits are just as obvious. These changes would allow for the analysis of how terrorist organizations may co-evolve to occupy particular “niches”, how terrorist organizations would compete for recruits, and how and when organizations might split (or even merge).

However, we can still get a great deal of mileage out of the simple single-terrorist-organization model. Until all of this model’s implications are fully understood, we should refrain from complicating it by making this change.

## **Splinter Cells**

### **Relevant Classes: Terrorist Brain, Terrorist Behavior Set**

One minor assumption that is made in the model and may be challenged is that isolated terrorist actors who become disconnected from their terrorist organization simply retreat back into the general population.

Alternatives to this assumption abound. As discussed above in “Only One Terrorist Organization”, these independent actors could act as the “seedlings” of new organizations.

These actors could also enter a “sleeper” mode, and might be provoked into acting if they become aware of other terrorists’ actions. (This follows Sun-Ki Chai’s analysis that ter-



rorist action is a form of inter-cell communication.) By acting independently, they still manage to boost estimations of terrorist competence, improve recruiting prospects, and perhaps even rebuild linkages to any remnants of the main terrorist organization.

Whatever choice is made with regard to handling isolated terrorist actors, we should be ready to defend our treatment of the situation.

I hope the preceding analysis proves helpful both for purposes of fine-tuning the terrorism model and for defending it against critical attack. In retrospect this is a “first-order” analysis that considers the model as formalized in Java code. I have not questioned the fundamental mechanics of the abstract model that is the basis for the Java formalization. As a result, all of the suggestions flowing from this analysis take the form of “if you want to improve or change X, modify code Y”. None of the suggestions involve reassessing the basic model itself – they are merely tweaks or add-ons. As a consequence, they all offer improvements at the expense of added complexity.

It might be helpful to follow-up with a deeper “second-order” analysis that looks past the code to the abstract model upon which the code is built. The implications of such an analysis may result in even more radical code changes, but it is entirely possible these changes could produce a cleaner, more elegant and even simpler model and Java formalization.



## References

---

- Ackerman, Gary, Bhattacharjee, Anjali, Klag, Matthew, Mitchell, Jennifer, 2002, "Literature Review of Existing Terrorist Behavior Modeling," Final Report to the Defense Threat Reduction Agency, 14 August 2002.
- Buckley, Walter, 1998, *Society - A Complex Adaptive System: Essays in Social Theory (International Studies in Global Change)*, Routledge.
- Chai, Sun-Ki, 1993, "An Organizational Economics Theory of Antigovernment Violence," *Comparative Politics*, Volume 26, Number 1.
- Cowan, George A., Pines, David, Meltzer, David, eds., 1994, *Complexity: Metaphors, Models, and Reality*, Reading, Massachusetts: Addison-Wesley Publishing Company.
- Dennett, Daniel C., 1991, *Consciousness Explained*, Boston, New York, Toronto, London: Little, Brown and Company.
- Dennett, Daniel C., 1995, *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, New York: Touchstone (1996 edition).
- Dennett, Daniel C., 2003, *Freedom Evolves*, New York: Viking.
- Eakin, Emily, 2003, "I Feel, Therefore I Am," *The New York Times*, April 19.
- Epstein, Joshua H., Axtell, Robert, 1996, *Growing Artificial Societies: Social Science from the Bottom Up*, Washington, D.C.: Brookings Institution Press.
- Jervis, Robert, 1997, *System Effects: Complexity in Political and Social Life*, Princeton, New Jersey: Princeton University Press.
- Hanle, Donald J., 1989, *Terrorism: The Newest Face of Warfare*, Washington: Pergamon-Brassey.
- Hanneman, Robert A., 1988, *Computer-Assisted Theory Building: Modeling Dynamic Social Systems*, Newbury Park, California: Sage Publications.
- Harmon, Christopher C., 2000, *Terrorism Today*, London; Portland, Oregon: Frank Cass.
- Holland, John, H., 1995, *Hidden Order: How Adaptation Builds Complexity*, Reading, MA: Addison-Wesley Publishing Company.
- Kauffman, Stuart A., 2000, *Investigations*, New York: Oxford University Press.
- Krebs, Valdis, 2002, "An Introduction to Social Network Analysis," <http://www.orgnet.com/sna.html>.
- Lanchester, R.W., 1956, "Mathematics in Warfare," *The World of Mathematics*, Vol. 4, Newman, J.R., ed., New York: Simon and Schuster.

- Michael K. Lauren, Michael K., Roger T. Stephen, Roger T., 2002, "Map-Aware Non Uniform Automata (Mana)—A New Zealand Approach To Scenario Modelling," *Journal Of Battlefield Technology*, Vol 5, No 1.
- MacArthur, Matthew Caleb, 2002, "An Adaptive Agent-Based Model of Dissent and Rebellion (Draft)," April 15.
- Newell, A., 1990, *Unified Theories of Cognition*, Cambridge, Massachusetts: Harvard University Press.
- Pape, Robert A., 2003, "The Strategic Logic of Suicide Terrorism," *American Political Science Review*, Vol. 97, No.3.
- Pigliucci, Massimo, 2000, "Chaos and Complexity: Should We Be Skeptical," *Skeptic*, Vol. 8, No.3.
- Reynolds, Craig W., 1987, "Flocks, Herds, and Schools: A Distributed Behavioral Model," *Computer Graphics*, 21(4) (SIGGRAPH '87 Conference Proceedings).
- Schelling, T. C., 1971, "Dynamic Models of Segregation," *Journal of Mathematical Sociology*, 1:143-86.
- Skyrms, Brian, 1996, *Evolution of the Social Contract*, New York: Cambridge University Press.
- Simon, Herbert A., 1981, *The Sciences of the Artificial*, 2nd. Edition, Cambridge, Massachusetts: The MIT Press.
- Weiner, Johnathan, 1994, *The Beak of the Finch: A Story of Evolution in Our Time*, New York: Vintage Books.
- Wolfram, Stephen, 2002, *A New Kind of Science*, Champaign, Illinois: Wolfram Media, Inc.

## **Acronyms and Abbreviations**

---

ABIR	Agent-Based Identity Repertoire
ABM	Agent-Based Modeling
CAS	Complex Adaptive Systems
CNA	Center for Naval Analyses
CRP	Central Research Project
DoD	Department of Defense
FBI	Federal Bureau of Investigation
GUI	Graphical User Interface
IDA	Institute for Defense Analyses
IKB	Integrated Knowledge Base
ITSD	Information Technology and Systems Division
MANA	Map Aware Non-uniform Automata
PSYOP	Psychological Operations
SNA	Social Network Analysis
SO/LIC	Special Operations/Low Intensity Conflicts
US	United States (of America)
V&V	Verification and Validation



REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. <b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b>					
1. REPORT DATE (DD-MM-YY) October 2006		2. REPORT TYPE Study		3. DATES COVERED (From – To)	
4. TITLE AND SUBTITLE Analyzing Adversaries as Complex Adaptive Systems				5a. CONTRACT NUMBER DASW01-04-C-0003 W74V8H-05-C-0042	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBERS	
6. AUTHOR(S) Dale E. Lichtblau, Task Leader, Brian A. Haugh, Gregory N. Larsen, Terry Mayfield				5d. PROJECT NUMBER	
				5e. TASK NUMBER IDA Central Research Program C5055	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESSES Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				8. PERFORMING ORGANIZATION REPORT NUMBER  IDA Paper P-3868	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				10. SPONSOR'S / MONITOR'S ACRONYM IDA	
				11. SPONSOR'S / MONITOR'S REPORT NUMBER(S) IDA Paper P-3868	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release, unlimited distribution: 5 January 2007					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT  The objective of this study was to assess information technology tools to counter asymmetric threats when considered as complex adaptive systems (CASs). We focused primarily on the use of agent-based modeling and simulation technology. This report describes both agent-based modeling (ABM) and an agent-based model developed to explore the utility of ABM technology to counter asymmetric threats. We conclude that while terrorist groups considered qua systems are undoubtedly adaptive, it is not obvious that they are complex in the strict theoretical sense of that term. As a consequence, it is not clear that terrorist threats are amenable to the analytic techniques afforded by complexity theory. Moreover, while ABM technology may offer significant value in many fields, it is not at all clear that the technology offers tactical value to counter these growing asymmetric threats. We argue that human behavior is too complex and too poorly understood to be accurately modeled in anything but a simplified and unenlightening way using the technology and agent-based modeling techniques currently available—particularly for tactical advantage. The real value—potentially inestimable value—lies in the systematic and methodical process of making explicit the assumptions regarding the fundamental factors governing agent behavior used in the models.					
15. SUBJECT TERMS Complex Adaptive Systems (CASs), Agent-Based Modeling (ABM), Complexity Theory, Counter-Terrorism					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Dale E. Lichtblau
Unclassified	Unclassified	Unclassified	Unlimited	97	19b. TELEPHONE NUMBER (Include Area Code) (703) 845-6683